Fall 2002

Gridpoints

The Quarterly Publication of the NASA
Advanced Supercomputing Division

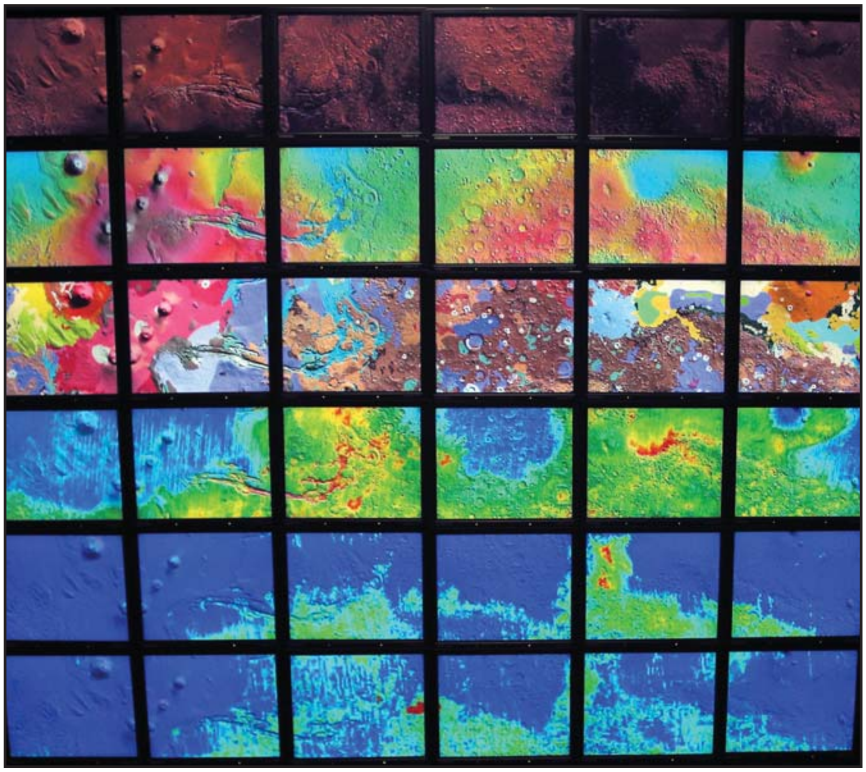

NAS researchers have developed an innovative visualization system, the 'hyperwall,' which allows scientists to view complex datasets and multiple parameters. See page 11

FEATURES

Features

4

NASA's IPG Team Develops New Tools

NASA's IPG team meets another of its milestones, creating tools to simplify job submission to distributed grid resources.

Holly A. Amundson

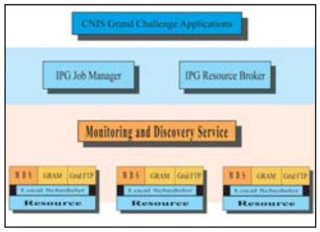

8

Nanotechnology: Simulating the Future of Computing

NAS scientists demonstrate how the laws of quantum mechanics affect transistor and device miniaturization.

Julie Jervis

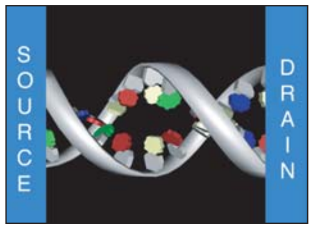

1A

NASA At SC2002

A look at NASA exhibits, events, and speakers at the annual high-performance networking and computing conference, SC2002, in Baltimore, Maryland, November 17-21, 2002.

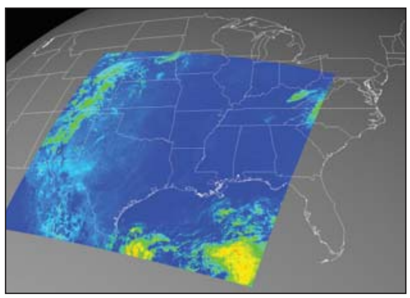

11

The Vision Behind the 'hyperwall'

NAS researchers have developed an innovative visualization system that allows scientists to view complex datasets and multiple parameters.

Julie Jervis

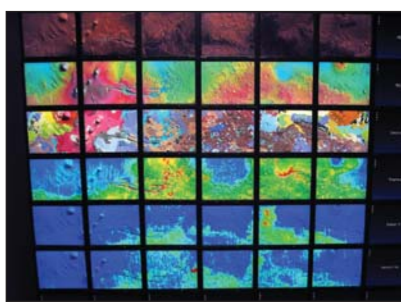

Also in this issue:

On The Cover:

Mars terrain data displayed on the NASA Ames-developed 'hyperwall" – a seven-by-seven cluster of flat panel screens, each driven by its own dual-CPU computer with high-end graphics card. Each row in this visualization projected on the hyperwall depicts the same 48-degree swath along Mars' equator from 180 degrees West to 180 degrees East. See page 11 for details.

(Dataset courtesy MOLA science team ; visualization by Glenn Deardorff)

Gridpoints
The Quarterly Publication of the NASA Advanced Supercomputing Division
Vol. 3, No. 3 • Fall 2002

Executive Editor: John Ziebarth

Editor: Nicholas A. Veronico

Writers: Holly A. Amundson, Jill A. Dunbar, Julie Jervis, Rob Van Der Wijngaart

Graphic Designer: Sue Kim

Illustrator: Cliff Williams

Photographer: Michael Boswell

Contact the Editor:
Nicholas A. Veronico
Gridpoints
M/S 258-6
Moffett Field, CA 94035-1000
Tel. (650) 604-3046 • Fax (650) 604-4377
nveronico@mail.arc.nasa.gov

The staff of *Gridpoints* would like to thank the following people for their assistance in preparing this publication: Michael Aftosmis, M.P. Anantram, Asen Asenov, Marsha Berger, Glenn Deardorff, Daniel Dorney, Reynaldo Gomez, TR Govindan, Chris Henze, Thomas Hinke, Vee Hirsch, Marjorie Johnson, William Johnston, Randy Kaemmerer, Creon Levit, Anthony Lisotta, Nagi Nicholas Mansour, John Melton, Parviz Moin, George Myers, Chuck Niggley, Marcia Redmond, Tim Sandstrom, Karim Shariff, Warren Smith, Judith Utley, Rob Van Der Wijngaart, Arsi Vaziri, Anne Woods, Jerry Yan, Pamela Walatka

Gridpoints is a quarterly publication of the NASA Advanced Supercomputing Division at NASA Ames Research Center. The division's staff includes government employees and contractors from Advanced Management Technology Inc., Computer Sciences Corp., Cray Inc., Eloret, Raytheon Co., and SGI.

Subscribe to *Gridpoints*

Subscriptions to *Gridpoints* are complimentary. Please send e-mail with your name, company name, title, address, and zip code to:
gridpoints@nas.nasa.gov
or subscribe online at:
www.nas.nasa.gov/gridpoints

Note: Subscribers outside the United States can access *Gridpoints* on the web in PDF format at: http://www.nas.nasa.gov/gridpoints

News From NAS

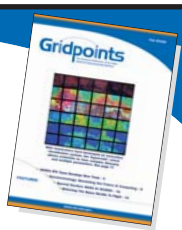

NAS Scientist Receives International Award

NAS senior scientist Deepak Srivastava recently received the Eric Reissner Medal for distinguished work in computational nanosciences. The medal, awarded in August at the International Conference on Computational & Experimental Engineering and Sciences (ICES), cites Srivastava's "distinguished work in computational nanosciences and carbon nanotubes."

The award, given by an international forum of scientists and engineers, recognizes Srivastava's contributions in broadening the understanding of the nanomechanics and electronic behavior of carbon nanotubes. It also signals an international acceptance of nanotechnology as the next growth area in the broader scientific and engineering communities.

Srivastava is technical lead of the NAS Division's computational nanotechnology investigations, and has concentrated his work on computational nanotechnology in nanomaterials, nanoelectronics, and nanodevices. Possible applications include ultra-strong, light-weight composite materials and molecular electronics devices for computing and sensing purposes, as well as concepts and materials for solid-state quantum computers.

Srivastava's new focus area will include state-of-the art implementation of

Continued on page 2

NAS Mission

To lead the country in the research, development, and delivery of revolutionary, high-end computing services and technologies, such as applications and algorithms, tools, system software, and hardware to facilitate NASA mission success.

From The Division Chief

For fifteen years, the NAS Division has taken a leadership role in support of the annual supercomputing conference, this year known as SC2002. Many of the division's research highlights shown in this issue of *Gridpoints* will be displayed at SC2002. (Details begin on page 1A.)

I am proud to announce that the division's researchers were recently honored with two prestigious awards. The Cart3D development team was the recipient of NASA's Software of the Year Award (see page 3), and the division's computational fluid dynamics research on the DeBakey Ventricular Assist Device (VAD) was recognized as NASA's Invention of the Year for 2001, announced in spring 2002 (see *Gridpoints*, Summer 2002, page 3). This award-winning work by our research teams is the culmination of more than a dozen years of computational science research and development for both products.

In addition to the Cart3D and DeBakey VAD highlights, division researchers will present a number of new tools and applications at SC2002 that form the infrastructure of NASA's Information Power Grid (IPG). The IPG team recently demonstrated the CODE (Control and Observation in Distributed Environments) framework, which enables researchers to control multiple resources, services, and applications in the grid environment (see page 4). To support the CODE framework, the IPG team also developed a grid abstraction tool that presents information from the MDS in a uniform manner.

Furthering our development of a robust infrastructure for the IPG, the division has released new versions of the NAS Parallel Benchmarks. GridNPB3.0 consists of four sets of problems that represent applications on computational grids for measuring grid performance (see page 14).

NAS researchers recently received an Information Sciences and Technology Directorate grant to develop the "hyperwall," a problem-solving environment consisting of a seven-by-seven cluster of flat-panel screens, each driven by its own dual-CPU computer. The hyperwall will help scientists use their own perceptive skills to analyze data by recognizing patterns, detecting features and changes, and sorting relationships between datasets.

NAS scientists are modeling the characteristics of electron flow in a wide variety of nanostructures, in an effort to demonstrate how the laws of quantum mechanics affect transistor miniaturization. These laws will aid researchers in building nanoscale transistors with an eye toward developing smaller-scale computers for future NASA space missions (see page 8).

From the need to solve fluid flows within the VAD, and during the development cycle of Cart3D, many important collaborations were established with researchers at other NASA centers, commercial companies, and numerous academic institutions. SC2002 provides our scientists with a venue to display their research while seeking new collaborators to help further NASA's mission goals.

I hope you enjoy this issue, and, as always, I welcome your feedback.

John Ziebarth
jziebarth@mail.arc.nasa.gov

Continued from page 1

simulation algorithms on NAS supercomputers, in a user friendly way. Srivastava hopes this area of research will pave the way for NAS to become a service provider for simulations in nanomaterials and molecular electronic devices and systems to other NASA centers, and to industry and university customers. NASA Ames' nanotechnology program plays a significant role in steering nanotechnology developments both within and outside NASA.

The Eric Reissner Medal is awarded every two years by ICES for excellence in computational mechanics research. Reissner, who passed away in 1996, was a noted teacher, researcher, and author, who expanded the foundations of the theory of mechanics, leading to advances in the design of both civil engineering and aerospace structures.

Biologists, Nanotechnologists, and Information Technologists Gather

Nearly 100 experts gathered at NASA Ames Research Center October 7-9 for the Biology, Information Science, Nanotechnology Fusion and NASA Missions Invitational Workshop to explore cross-disciplinary approaches to solving the science and technology challenges facing future NASA missions. Organized by Ames Research Center and the Universities Space Research Association (USRA), the workshop was the first of its kind, designed to explore the emerging fusion of these three fields of science.

"Our goal was to start a conversation and gather new ideas," says T.R. Govindan, NASA Advanced Supercomputing Division senior scientist and co-chair of the workshop. "This is a very diverse group of people, and everyone speaks their own language, but the presentations gave us excellent examples of what this research might accomplish for NASA missions."

Experts from universities, industry, and government agencies throughout the United States joined NASA personnel for three days of presentations, panel talks, and discussions covering topics that included the use of biology to solve information technology challenges, life detection and sensors, computational biology and bioinformatics, and nanotechnology assembly.

By inviting mission experts from all NASA enterprises to talk about their needs and challenges, the workshop leaders began the task of matching these nascent research areas to NASA's mission needs.

Participants Baruch Blumberg, director of the NASA Astrobiology Institute, and Michael Storrie-Lombardi of NASA's Jet Propulsion Lab both agreed the meeting was an important first step in bringing these sciences together, and hoped it would lead to a concentrated effort to follow up on the leads and themes discussed. Govindan says the next step will be to publish a report detailing the themes, ideas, and suggestions that emerged from the workshop, and provide a framework for planning. "Based on the level of enthusiasm and interest in the fusion of these three fields, it's likely that this kind of conference might become a regular event," he says. For more information visit: http://binfusion.arc.nasa.gov

Representative Honda Visits Ames' Aviation Security Laboratory

Representative Mike Honda (second from right), of the 15th Congressional District in Silicon Valley, toured the NASA Advanced Supercomputing (NAS) Facility at Ames Research Center on August 23. The facility houses NASA Ames' new Aviation Security Laboratory, a mock airport setting designed to investigate where information technology can be applied to the flow of passengers through an airport, to detect potential security concerns. Aviation Systems Division Chief Tom Edwards, (far right), and Thomas Hinke, lead scientist for the Ames Aviation Security Lab (second from left), look on as Bill Thigpen, NAS Engineering Branch chief, explains how the system works. Honda serves on the House Budget Committee, the House Committee on Science, including its subcommittee on Research, and on the Transportation and Infrastructure Committee, including its subcommittees on Aviation, Highways and Transit, and Water Resources and Environment. (NASA/Tom Trower)

Recognizing Leadership in Scientific Computing

Every year NASA recognizes outstanding accomplishments of the agency's research community by presenting the NASA Software of the Year award. This year, the prestigious award was placed in the hands of the NASA Advanced Supercomputing (NAS) Division's Michael Aftosmis, John Melton of NASA's Aerospace Directorate, and Marsha Berger of the Courant Institute, New York University. The team was recognized for their long-time work on Cart3D, a design and analysis software package for complex aerospace vehicles (see *Gridpoints*, Summer 2000, page 4).

Co-sponsored by NASA's Chief Information Officer and NASA's Inventions and Contributions Board, the Software of the Year award is judged by the NASA Software Advisory Council. To meet eligibility requirements, software must be supported, adopted, sponsored, or used by NASA, it must be significant to the NASA mission, and be of commercial grade. In addition to Cart3D, a Direct Simulation Monte Carlo software package, called DAC, from NASA's Johnson Space Center, Houston, was named co-winner of the annual award.

What sets Cart3D apart from other Cartesian mesh generation packages is its ability to accept geometry of arbitrary complexity and process it in a fully automated manner. Reynaldo Gomez, the Space Shuttle aerodynamics subsystem manager at NASA Johnson is especially pleased with the robustness of Cart3D: "The real power of Cart3D lies in its simplicity of use. Historically, computational fluid dynamics (CFD) codes required a lot of manual intervention and training to develop grid systems and to run the flow solver. With Cart3D, I've actually taken a third-year Air Force cadet who had no background in CFD, no knowledge of UNIX or CAD systems, and within a period of about a week, he was able to generate grids, start running the flow solver, and begin postprocessing results for Space Launch Initiative configurations."

In addition to its applications in NASA programs such as the Space Launch Initiative, Cart3D has significant relevance in the commercial world. In September 2001, ICEM CFD Engineering, Inc. of Berkeley, Calif., was granted an exclusive license by NASA Ames' Commercial Technology Office in the areas of aerospace, automotive engineering, and turbomachinery. Since then, the company has also been granted a non-exclusive license in the areas of electronics, electromagnetics, and the process industry – injection molding and casting.

Another company based in Seattle, Washington, also has a commercial license for Cart3D, and is using the grid generation software to investigate the American Airlines Airbus crash on November 2001 on Long Island, New York. "This kind of 'virtual flight' capability allows people to do aero forensics on this type of crash," explains Aftosmis. The Airbus crash analysis includes estimations of aerodynamic loads imposed by turbulence from other aircraft in the area.

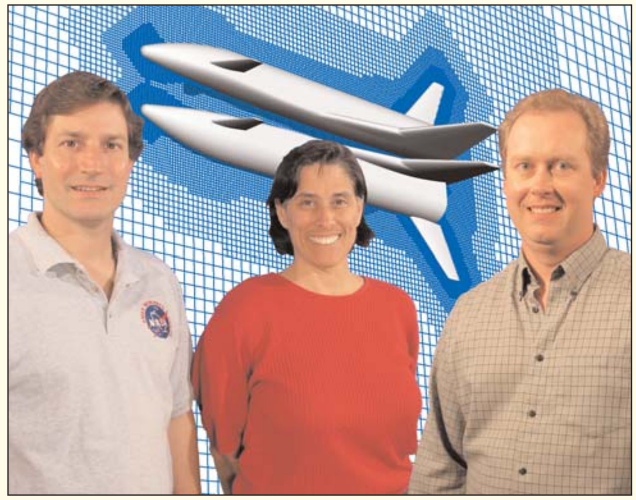

NASA researchers Michael Aftosmis (left) and John Melton (right) and Marsha Berger, from New York University's Courant Institute, co-developed Cart3D

(Michael Boswell)

Cart3D's ease of use and automation makes it possible to carry out virtual flight, where every permutation of a design is explored over the complete range of expected flight conditions. "Cart3D provides designers with the trend information they need in order to understand how their designs are going to operate, not just under a single condition, but over a wide range of flight conditions," explains Melton. "Cart3D makes the computational fluid dynamics process fast enough to replace lower-order design methods with high-fidelity techniques – the key to streamlining the design and analysis process of complex aerospace vehicles," adds Aftosmis.

The level of automation demonstrated by Cart3D has sparked international interest in Cartesian mesh generation. "It is very gratifying to see this line of research being picked up by other researchers worldwide," says Berger. The team hopes to continue generating interest in Cartesian meshes as they add new capabilities to Cart3D.

For more information about the Cart3D software package, visit: www.nas.nasa.gov/~aftosmis

NASA's IPG Team Develops New Tools

NASA's IPG team meets another of its milestones, creating tools to simplify job submission to distributed grid resources.

NASA's Information Power Grid (IPG) team has taken another step toward making the agency's prototype grid environment more user-friendly. To support the Computing, Information, and Communications Technology (CICT) Program's Computing, Networking, and Information Systems (CNIS) Project in reaching one of its major milestones in September, the IPG team delivered a grid programming environment consisting of the Globus Toolkit and tools developed by the IPG team. Coincidentally, the tools developed for the CNIS milestone enabled the IPG team to meet one of its own goals. In addition, the IPG team's development work provided access to as many as 17 grid resources at seven locations. The heterogeneous computer resources included Sun, Linux, Linux cluster, and SGI machines running Irix, Solaris, and Linux operating systems. CICT researchers used 13 resources at four locations to meet the milestone – 10 resources in four locations was the measure of the milestone's success.

In achieving this goal, IPG team members at NASA Ames, Glenn, and Langley Research Centers collaborated with researchers at the National Center for Supercomputing Applications (NCSA) in Illinois, Information Sciences Institute (ISI) at the University of Southern California in Marina Del Rey, and Argonne National Laboratory (ANL) in Argonne, Illinois. Together, these organizations provided the support and expertise needed to make the resources available from multiple locations. To make it easier for users to take advantage of these geographically distributed resources, several new grid tools were developed and deployed. These tools use the Globus Toolkit, which provides a uniform interface to the resources.

The tools created for simplifying the process of submitting jobs to the distributed grid resources include a job manager to run specified jobs and track job progress, an abstraction tool for handling job scheduler information on multiple machines, and a resource broker for selecting grid resources (see figure, page 5). Standard grid security measures were also put in place on every resource, ensuring the secure transfer of all data.

Tools of the Trade

To oversee all the "transactions" taking place among the grid resources, the IPG team used an existing framework called Control and Observation in Distributed Environments, or CODE (see *Gridpoints*, Summer 2002, page 18). The CODE framework, created to provide a standardized process for making observations and performing actions on remote computer systems, served as the basic communication mechanism for the CICT milestone. "We had to accommodate Java and Perl clients, so we needed a communication framework that could do both – CODE was the answer," explains IPG team member Warren Smith, who developed the framework. CODE was also used to manage events – automatic notifications were sent to users alerting them upon completion of a task, or when a machine became unavailable.

The job manager delivered to IPG users is designed to submit, track, and cancel jobs. Users describe the set of files they want to stage (transfer) before they run their job, which application they want to execute, and where they want their files to end up when the jobs are complete. "The job manager is really tuned to be able to run jobs reliably, once users define what they want to accomplish and where they want the results to go," says Smith. To move files from one machine to another within the grid environment, the job manager takes advantage of the GridFTP mechanism, a method for transferring files. To execute applications on grid machines, the job manager utilizes the Globus Resource Allocation Manager (GRAM), a service that enables remote job submission.

The information abstraction tool, also developed for the milestone, is intended to take the "noise," or extra details, out of using grid resources. The abstraction tool comprises components that represent information about computers, networks, and schedulers, and encapsulates information from a hierarchy of servers called the Monitoring and Discovery Service (MDS). "The problem is that the format of the information in the MDS changes periodically, and this is a nuisance to users. The abstraction tool hides these changes, providing a consistent view of the information needed by users to run their jobs," explains Smith.

To handle resource selection within the grid environment, the IPG team also created a resource broker. To utilize the resource broker, users are only required to describe the job they want to run by specifying the number of CPUs needed, the amount of memory needed, the operating system requirements, and any other information unique to their jobs. Once this information is collected, the broker uses the description and data from the abstraction tool to determine which machines satisfy their job requirements. The broker then generates a list of machines capable of running the job, prioritized by current loads on the machines' processors. "Our application utilized the IPG resource broker software to select the right compute host for each computation, greatly simplifying our submission process," says NAS researcher Stuart Rogers.

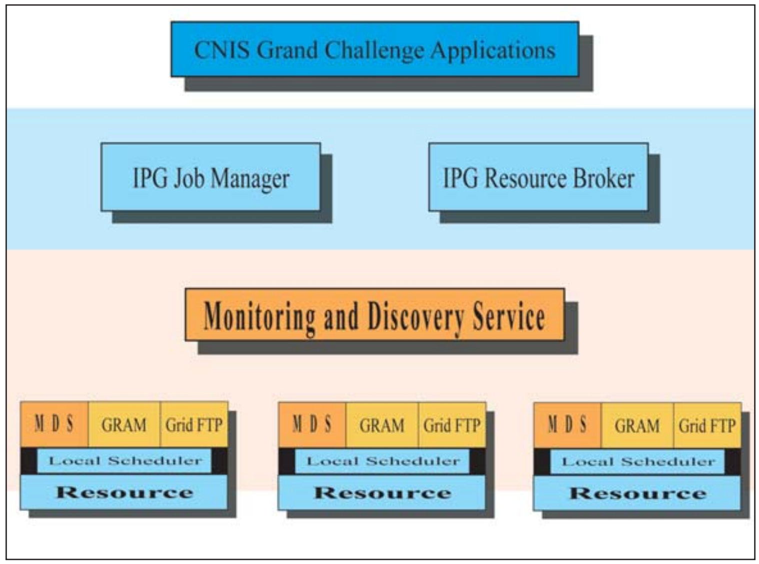

This diagram represents a conceptual view of the IPG software environment, including software components developed by the IPG team used to meet the requirements of a Computing, Information, and Communications Technology (CICT) Program milestone. The IPG resource broker obtains information from the various components of the Monitoring and Discovery Service, an information service containing resource information, in order to suggest where to run jobs. To execute applications on grid resources, the job manager utilizes the Globus Resource Allocation Manager, a service that enables secure remote job submission, and the GridFTP mechanism to securely transfer files between grid resources.

Lessons Learned

"This milestone provided us with a great opportunity to study the interoperability of grids, helping uncover some grid management problems. This will lead to future work in making grids more compatible with one another," explains NAS IPG task lead Tony Lisotta. One of the biggest challenges the IPG team faced while working to meet this milestone was ensuring the use of compatible grid software so the resources could "talk" to one another. Contributing to this challenge, the grid community was in the middle of a Globus upgrade. During this transitional period, a lot of backward compatibility issues surfaced – centers were using different versions of the Globus Toolkit, preventing resources from communicating with one another.

To promote communication about Globus Toolkit versions and other issues that came up during the process of meshing the various grid environments, several practices were put into place. E-mail lists were established so that everyone involved had access to information about problems and solutions; weekly meetings were held for researchers working on the applications portion of the milestone; and regular conference calls were conducted to encourage communication among all participants.

"Coordination and collaboration are some of the most critical challenges in managing a grid. Lessons learned with this milestone are being shared in the grid community as IPG takes a lead role in making grids work, and work well together," explains IPG team member Judith Utley. IPG team members and collaborators plan to write a "lessons learned" paper addressing grid management, which will be submitted to the Global Grid Forum (GGF) for discussion. In addition, a new group, dedicated to addressing grid management issues, and co-chaired by Utley, has been formed within GGF. "This will contribute to the overall grid community in the long run, as we identify areas where standards and agreements are necessary," adds IPG team member George Myers.

Promising Future for Grids

While not all the tools and practices used to meet this milestone will become standard grid practices, integrating some of these new tools into the IPG environment will make it more robust and increase its capabilities. "Given what we've learned about these tools, we are reviewing them to see how they should be changed, and what we can do with them. The job manager, for example, will be extended a bit to add some more functionality," says Smith. "The effort to meet the milestone led to the development of tools that will continue making the IPG easier to use – that has been our goal from the beginning, and will continue to be the theme as we develop the tools we envision," says Myers.

The September milestone provided the IPG team with an opportunity to demonstrate its ability to contribute to the advancement of grids. Says Myers: "I think it was an amazing accomplishment to bring together such an array of diverse organizations, and it says something about the power of grid technology that we could pull it all together."

— *Holly A. Amundson*

Editor's note: *The AeroDB group, part of the CICT/CNIS Project, used tools provided by the IPG team for this milestone to implement an application for automating large CFD parameter studies. This application enabled the AeroDB group to analyze a flight vehicle over a wide range of flight conditions. Look for an article outlining this portion of the milestone in the Winter 2003 issue of* Gridpoints.

Center for Turbulence Research Summer Program Making Worldwide Waves

Presentations and papers from the popular NASA-Stanford University program attract international attention.

After four weeks of intense collaborative study, participants in the Center for Turbulence Research (CTR) summer 2002 program presented their findings before a large international audience at Stanford University. The August presentation ceremony marked the final step before publication and worldwide distribution later this year.

Since its first session in 1987, CTR's biennial summer program has attracted research proposals from laboratories all over the world; each evaluated according to its relevance and match with NASA's interests. This year, the 46 selected participants came from 11 countries, nine states, and 40 institutions across the globe. Research focused on nine areas of CTR's expertise – turbulence acoustics, Reynolds-Averaged Navier-Stokes (RANS) modeling, large eddy simulation (LES), numerical methods for LES, turbulence fundamentals, stratified shear flow turbulence, flow optimization, nanofluidics and biology, and turbulent combustion and sprays.

"The summer program is a cost-effective way for researchers to get a first cut at their ideas, which are typically brand new, exploratory, and often high risk."

—Parviz Moin
Center for Turbulence Research

A cooperative program between Stanford University and NASA Ames Research Center, CTR serves as a consortium for turbulence research by bringing theorists and experimentalists together in a number of research programs. The NAS Division co-manages and supports CTR by providing supercomputing power and access to its huge archive of raw field data, enabling researchers to test their ideas. NAS scientists are also involved in several individual CTR research projects.

"The summer program is a cost-effective way for researchers to get a first cut at their ideas, which are typically brand new, exploratory, and often high risk," says Parviz Moin, CTR's director. Each of the 32 projects conducted this year were based on the researchers' own proposals, some of which used the full potential of NAS' supercomputing resources. "NAS had to increase CTR's allocation of computer time threefold to accommodate their use of visualization facilities, networks, and data archives," says Nagi Mansour, CTR's deputy director and lead of NAS' Physics Simulation and Modeling Office.

"Historically, turbulence researchers have been preoccupied with generating data to test ideas and take measurements," explains Moin. "The numerical simulation techniques pioneered at Ames provide raw field data, so researchers do not have to worry about producing the data. This is being emulated all over Europe, as it provides a single standard for independent testing."

This year's program started with registration on July 28, followed by an intense work schedule that included kick-off presentations, midterms, tutorials, and final presentations. The end-of-term presentations included an overview of methods, investigations, comparisons, and conclusions, while the research teams faced close scrutiny and intense questioning from experts in the audience. These experts included guests from other institutions and government agencies that might be interested in funding the participant to continue the research at their home institute.

Moin says that past findings from the program have had a major impact in the field of turbulence control and modeling, resulting in methodologies used for standard code throughout the world in the areas of simulation technology and database diagnostic tools. This year, the monthly journal, *Physics of Fluids,* published by the American Institute of Physics, will publish a summary of the findings.

Karim Shariff is a NAS research scientist who has been involved with the program from the start when he participated as a graduate student. Over the years, he has participated in several successful projects, including *Simulation and Modeling*

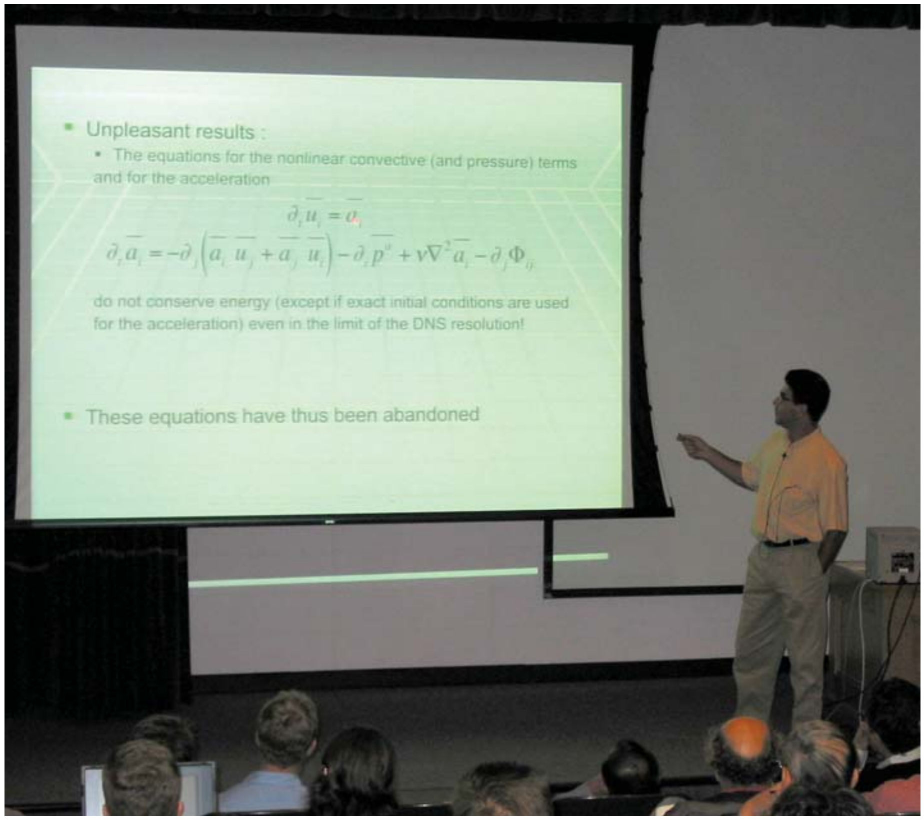

CTR Summer program participant Daniele Carati, of Brussels University in Belgium, presents his group's findings on Large Eddy Simulations. CTR invited 46 scholars to participate in the summer program. They represented 40 institutions from nine U.S. states and 11 countries.

of the Elliptic Streamline Flow in 1996, which was the first study of turbulence in this type of flow. This summer, Shariff was one of 26 hosts from Stanford and NASA, and guided a study on contrail formation in aircraft wakes.

"I've always been interested in aircraft trailing vortices," says Shariff. "Atmospheric scientists have suspected for some time that airplane contrails, those two lines that come out like smoke behind an aircraft, have an impact on climate. Basically, contrails are clouds of ice that form when soot or aerosols from engines cause condensation of water vapor present in the engine exhaust and in the atmosphere. Contrails artificially increase cloudiness and trigger the formation of cirrus clouds, which in turn alters the climate on local and global scales."

The proposal to study the conditions that cause the formation of contrails came from Roberto Paoli, a post-doctorate scientist with the Computational Fluid Group at CERFACS, the European Centre for Research and Advanced Training in Scientific Computation, in France. Paoli's research project simulated the interaction of jet exhaust with the vortex of the aircraft and tracked representative soot particles to show when ice begins to accumulate, what the radius of the ice is going to be, and how the distribution of ice particles evolves over time. Using prior knowledge gained by researchers in the Earth Science Division at Ames, who have been doing flight studies of cloud formation and contrails for many years, the scientists were able to identify the areas of the parameter space that are important.

After completing his calculations, Paoli concluded that vapor first supersaturates at the edges of the engine jet exhaust, then the vortex increases the mixing of water vapor and air, which favors condensation. Paoli also concluded that as ice particles form, they decrease the amount of vapor in the atmosphere, resulting in a significant deviation from the classical theory of a mixing line. These kinds of results will ultimately be used by Earth scientists as models in simulations of the global climate.

Like all the participants in CTR's summer program, Shariff hopes to see more research coming out of this new and innovative project. Among the many offshoots from this summer's work, Paoli is planning further study into how jet mixing is suppressed within the vortices, a study that Shariff will find useful for his own work on the radar detection of aircraft trailing vortices (see *Gridpoints*, Fall 2001, page 14).

"This is really the first time this type of simulation has been done," says Shariff. "If it were not for the summer program, Paoli would be working on it in Europe, and I might never hear about it. This is a wonderful opportunity for people doing cutting-edge research in different but related fields to come together, bring each other up to speed, and benefit from sharing valuable experience and knowledge."

— *Julie Jervis*

For more information on the Center for Turbulence Research, visit:

http://ctr.stanford.edu

Shariff's work on aircraft wake vortex detection can be viewed at:

http://people.nas.nasa.gov/~shariff

Nanotechnology: Simulating the Future of Computing

NAS scientists demonstrate how the laws of quantum mechanics affect transistor and device miniaturization.

Since transistors were invented in 1947, their miniaturization and large-scale integration have been so rapid, that chip designers have been able to double computer processing power approximately every 18 months. This rule-of-thumb is commonly known as Moore's Law, after Gordon Moore first made this observation in 1965. Today, microprocessors the size of a fingernail contain tens of millions of transistors and are used in supercomputers with enough processing power to simulate complex models of Earth's climate system.

At the same time, developing expertise for modeling and understanding the characteristics of electron flow in a wide variety of nanostructures is a goal that NAS scientists hope will one day enable them to send a supercomputer the size of a briefcase into space. With such sophisticated computing power, a spacecraft would be more autonomous, be able to self-diagnose, self-repair, and better guide mission control on Earth into making the right decisions in outer space.

As transistor miniaturization has led to faster and more powerful chips, chip designers are now facing new challenges from the laws of physics. In its *2001 International Technology Roadmap for Semiconductors*, the Semiconductor Industry Association predicted that transistors will shrink to tens of nanometers within the next decade (see *How Small Is Nanoscale?)*. This raises concern among researchers and industry that, as the laws of quantum mechanics take over, current methods of lithographic chip etching will be inadequate and devices may not work as expected. Quantum mechanics is the science that deals with the behavior of matter and light at the atomic and subatomic scales. This dictates a different realm of physics than the one we are used to in the real world, causing tiny transistors to behave in a completely different manner than their conventional, larger siblings. Industry is turning to the research community for critical studies into new analytical modeling techniques, and to ensure a full understanding of the physics and chemistry involved at the quantum level.

"Part of our focus is to downscale semiconductor technology and model the current voltage characteristics of transistors with channel lengths less than 25 nanometers in size," says NAS research scientist M.P. Anantram. "We can simulate these transistors on the computer, determine the ideal design parameter space, and predict device characteristics. This can then guide experimentalists to build these transistors in a reproducible manner."

Together with colleagues Alexei Svizhenko and NAS Senior Scientist T. R. Govindan, Anantram is harnessing the power of NAS' parallel supercomputers to conduct the complex and highly sophisticated calculations necessary for nanoscale modeling. The process involves running thousands of calculations on potential transistor models, solving large sets of equations, and then using the results to continually refine the model (see figure on page 10). Just to compute the current voltage characteristics of one device can take as much as ten hours of computing time on 64 processors of an SGI Origin 3800 supercomputer.

NAS' supercomputing power has been instrumental in attracting international collaboration and advancing research in the field. Asen Asenov, head of the Department of Electrical Engineering at the University of Glasgow in Scotland is one of NAS' collaborators in the field of nanotransistors. "We chose NASA Ames for collaboration because they have, in my view, the most sophisticated and accurate quan-

How Small Is Nanoscale?

- One inch is equal to 25.4 million nanometers.
- The average width of a human hair measures 100,000 nanometers.
- One human red blood cell measures 10,000 nanometers wide.
- A single streptococcus (strep throat) group A bacteria measures 600 to 1,000 nanometers in diameter.
- A single HIV virion measures 80-100 nanometers in diameter.

tum transport simulator in the world, and the necessary computing power to run meaningful design-oriented simulations," says Asenov.

Asenov's team specializes in developing 3-D statistical simulations to understand fluctuations in the electrical characteristics of nanoscale devices. For each device, the Glasgow scientists use the results from NAS scientists' advanced 2-D quantum mechanical simulations to determine the appropriate parameters to use in their less sophisticated but faster 3-D atomistic quantum device simulator.

Asenov is planning to introduce more of the advanced quantum simulation methods developed by NASA Ames into the university's simulation tools, and to continue collaboration, especially in developing new, close to first-principle simulation techniques to describe the complex transport properties of atomic-scale devices.

Another successful aspect of Anantram, Govindan, and Svizhenko's work has been examining many of the physics issues concerning the flow of electrons in atomic-scale carbon nanotubes. Building devices using carbon nanotubes is a new, alternative approach that assembles larger objects from molecules, rather than starting with something large and trying to define smaller and smaller patterns. "Carbon atoms form a long tubular structure, which can either be a metal or a semiconductor, depending on the chirality, or 'twist,' of the tube," explains Anantram. Copper wire is not atomically precise, and therefore subject to defects, which cause scattering and reduces its current-carrying capacity. Carbon nanotubes *are* atomically precise, and the NAS researchers are exploring issues such as the ultimate current-carrying capacity, coupling to metal contacts and their electromechanical properties. "Chirality is a degree of freedom that nature offers to us, and we are interested in studying the best chiral angles for various nano-device applications," explains Anantram.

More recently, the NAS researchers have modeled the electromechanical properties with nano-sensor applications in mind. These investigations have yielded significant results in the scientific community. Working in collaboration with Amitesh Maiti of Accelrys Inc., a San Diego-based provider of simulation and informatics software and services, Anantram and Svizhenko were able to offer an explanation for an experiment published in the journal *Nature* in June 2000. In the paper, "Reversible Electromechanical Characteristics of Carbon Nanotubes under Local Probe Manipulation," conducted at Stanford University, scientists set a carbon nanotube between two metal contacts and pressed down on the middle of the tube using an Atomic Force Microscope (AFM) tip. They were surprised to find that this strain caused the current to decrease by a factor of 100, and theorized that the nanotube becomes a bad conductor of electricity because the bond structure changes in the area where the tip strikes the tube.

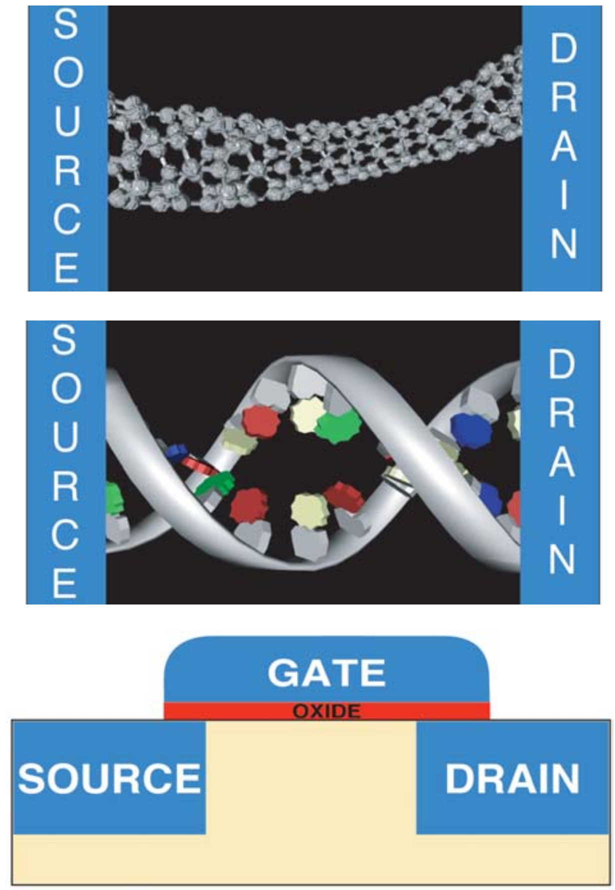

NAS scientists have developed models to demonstrate how the laws of quantum mechanics affect transistor and device miniaturization. The top model depicts a nanotube between source and drain contacts, where its ultimate current-carrying capacity, metal-nanotube coupling, and electromechanical/nano-sensor properties are being studied. The center model is a DNA strand used in transport experiments where the physics of conduction is being explored. At bottom is a nano-transistor, where 2-D doping profiles, polysilicon depletion, source-to-drain and gate tunneling, and the role of scattering are being studied.
(M.P. Anantram/Cliff Williams)

"We wanted to find out why this happened in the experiment, and predict what would happen if the experimentalists used nanotubes of other chiralities and diameters of nanotubes," says Anantram. The team began by modeling a nanotube made from a few thousand carbon atoms, simulating the effect of the AFM tip, and then examining the new configuration of atoms. They were particularly interested in seeing how this effect varied from just bending a nanotube. "What we found was that stretching was more important than bending," says Anantram. "The change in conductance is actually due to the 'rubberband effect' – strain across the entire tube – and not just because of the deformation in the central region by the AFM tip, as previously theorized."

The results of the team's collaborative study were published in a paper titled, "Electronic Transport through Carbon

Nanotubes: Effects of Structural Deformation and Tube Chirality" in the American Physics Society's journal, *Physical Review Letters*, March 25, 2002. This study continues to be the focus of ongoing research into the metallic and semiconductor properties of carbon nanotubes.

A third and more controversial paradigm for device design is the nascent research into conductivity in DNA. NAS scientists are interested in this form of nanotechnology for two main reasons: Currently, experimentalists around the world are getting contradictory results from their studies, providing a major challenge in establishing the true nature of DNA. In addition, they are interested in learning whether the arbitrary sequences that can be reproducibly designed by biochemists can be exploited in nanotechnology.

DNA has four building blocks: A, T, C, and G, which biochemists can arrange within the DNA structure in any arbitrary and reproducible manner – a huge attraction for nanotechnologists. The challenge has been isolating individual strands from the millions that are produced at one time, and locating them in an experiment. However, experimentalists have recently been able to take strands approximately 8-100 nanometers long and place them between metal electrodes to measure the current as a function of voltage, as though it were a nanodevice.

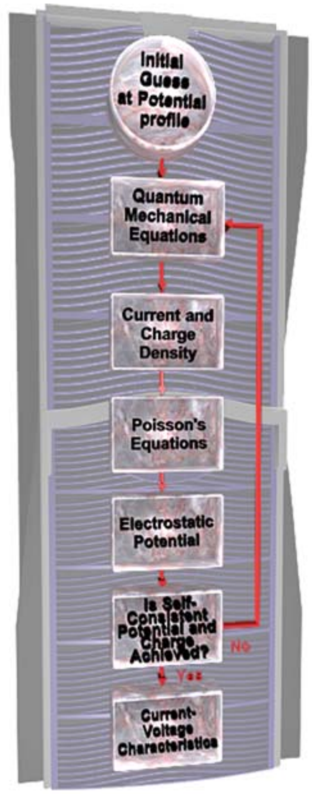

Scientists use this process for running quantum mechanical calculations of the current-voltage characteristics of a potential transistor profile. (Cliff Williams)

"These experiments have been very controversial," says Anantram. "One group in the Netherlands finds the strands are semiconducting in nature, and another in Switzerland finds they are metallic. Their experimental techniques differ and they have slightly different ways of doing these measurements, so the challenge is for us theory and modeling researchers to find out what is really happening."

So far, Anantram says, calculations performed in collaboration with NAS postdoctoral fellow Christophe Adessi (now a faculty member at the University of Lyon, France), and Steve Walch from the Center for Nanotechnology at NASA Ames, indicate DNA is, at best, a semiconductor. In addition, as the length of the DNA increases, it becomes an extremely poor conductor of electricity. The resistance increases much more rapidly than linearly with length. Though the researchers have two papers in review explaining some aspects of the experimental work, Anantram says that the physics of conduction in DNA is far from resolved, and there is much to be understood – for example, how the current-voltage characteristics change as a function of DNA sequence, and can DNA design obtain specific device characteristics.

As an overlap in the fields of information science and biology continues to grow, NAS scientists are developing the theory and simulations that will provide a deeper understanding of what will work and why. "Large scale simulations give us a better understanding of these emerging fields," says Anantram. "Technology is interrelated, and nanomachines will help us make the next leap forward."

— Julie Jervis

Semiconductor Modeling

The heart of modeling electronic nanotransistor devices is the Metal Oxide Semiconductor Field Effect Transistor (MOSFET), commonly recognized as the workhorse of modern-day computing. The MOSFET works like a super-efficient switch with three terminals: source, drain, and gate (see models on page 9). The gate is an electrode attached to the base of the transistor with a piece of positively charged (p-type) semiconductor in the middle and a thin layer of an insulator, such as silicon dioxide. When low voltage is applied to the gate, electrons in the form of current are able to flow freely from the source to the drain. But applying high voltage to the gate has the same effect as standing on a flexible water hose – electrons are repelled and fewer can flow through (see *Gridpoints*, Summer 2000, pages 10-13). In conventional MOSFET design, technical issues such as current leakage and electron tunneling can be solved using classical physics equations, but in devices that are less than 100 nanometers in size, these equations no longer describe the physics of resistance. Anantram, Govindan, and Svizhenko's work has been to simulate charge flow in tiny MOSFETs using quantum mechanical equations, and to provide answers to complex technical issues without the billion-dollar cost of building machinery and fabrication plants. "Going forward, no one knows whether silicon will be the best material for nanotransistors," says Anantram. "Our modeling effort can provide answers and guidelines that could eventually help companies by focusing on technologies that have a high chance of success."

SC2002 High Performance Networking and Computing

A look at NASA exhibits,
events, and speakers at SC2002

A special section of Gridpoints Magazine, Fall 2002

Welcome, SC2002 participants.

I'm pleased to welcome you to a conference that represents everything NASA is all about: pioneering the future, pushing the envelope, and doing what has never been done before.

The city of Baltimore will provide a wonderful setting for productive professional interchange about the latest developments in information architectures, scalable computing, data analysis, applications, and collaborative technologies.

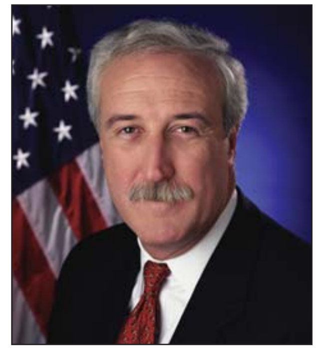

Bill Ingalls/NASA

Sean O'Keefe
NASA Administrator

For our part, NASA will demonstrate at SC2002 how we are utilizing advanced supercomputing technologies to improve aviation safety and efficiency, probe more deeply into the mysteries of the universe, model terrain and pinpoint possible spacecraft landing sites on Mars, and work to better understand the dynamics of Earth's climatic system.

We're also delighted to showcase how the adept use of computational science was a key factor in the development of the NASA-(Dr. Michael) DeBakey Ventricular Assist Device, NASA's Commercial Invention of the Year. NASA computational techniques that normally model rocket fuel flow helped to perfect this miniature heart assist pump, which promises to improve the lives of thousands of heart patients.

When we think about touching lives, clearly the kind of work all of us are engaged in has the potential to help inspire the next generation of explorers, which not so coincidentally is one of NASA's core mission goals. Given the fact that fewer young people are choosing to enter technical fields, I'm very pleased that this conference offers excellent programs that will help teachers apply the exciting work of computational research into classroom activities. By assisting teachers in their efforts to connect students to the technology frontiers you explore on a daily basis, I believe that SC2002 will help in immeasurable ways to build a better America.

—Sean O'Keefe

NASA is showcasing many of its exciting scientific research projects at SC2002. The following are descriptions of demonstrations, paper presentations, videos, and panel discussions listed by participating NASA centers and collaborators.

Ames Research Center

Mountain View, California

Ames Research Center (ARC) and its personnel work to develop technologies that enable the Information Age, expand the frontiers of knowledge for aeronautics and space, improve America's competitive position, and inspire future generations. Ames specializes in research geared toward creating new knowledge and new technologies that span the spectrum of NASA interests.

Atomistic Simulations on USA and Japan Grid

Researchers will demonstrate results from multiscale quantum-mechanical and/or classical atomistic simulations on a grid of geographically distributed PC clusters, to study environmental effects of water molecules on fractures in silicon.

The multiscale simulation approach seamlessly combines atomistic simulations based on the molecular dynamics (MD) method and quantum mechanical (QM) calculations based on the density functional theory. The multiscale MD/QM simulation code has been grid-enabled using a modular additive hybridization scheme; multiple QM clustering; and computation/communication overlapping. A preliminary code run has achieved a parallel efficiency of 94 percent on 25 PCs distributed over three PC-clusters in the United States and Japan. A video trip through a silicon crack tip shows water-silicon reactions.

Cart3D: A Package For Automated Grid Generation and Aerodynamic Database Creation

Cart3D, NASA's Software of the Year for 2002, is a software package for fully automated geometry processing and computational fluid dynamics (CFD) simulation. The software enables users with access to networked computing clusters or supercomputers to synthetically generate the types of aerodynamic databases typically obtained only through extensive wind tunnel testing. SC2002 attendees can select and manipulate pre-built vehicle geometries, and watch while meshes suitable for CFD analysis are automatically generated on demand. An entire suite of CFD runs for the chosen configuration can then be submitted for processing on NASA's Information Power Grid. These simulation results constitute a synthetically generated aerodynamic database that gives analysts information needed for design. Animations, images, and results from previously computed simulation suites will also be available for playback, discussion, and demonstration. Additional information is available on the web at: www.nas.nasa.gov/~aftosmis/cart3d

The Hyperwall: A Multiple Flat-Panel Display Wall For High-Dimensional Visualization

The NASA Ames "hyperwall" consists of a seven-by-seven matrix of flat panel displays driven by 49 rack-mounted dual-CPU nodes, each with its own high-end graphics card. The hyperwall emphasizes displaying multiple independent but related images, and provides useful means for composing and controlling these image sets. In place of elaborate software or hardware crossbar switches, researchers rely on the human visual system for integration, synthesis, and pattern discrimination in complex and high-dimensional data spaces.

The tabular layout of the hyperwall supports many existing "multiview" visualization paradigms, including spreadsheet-style approaches, multidimensional/multivariate techniques, and brushing/linking.

DeBakey Ventricular Assist Device NASA's Invention of the Year – 2001

The use of computational fluid dynamics (CFD) technology has led to several major design improvements on the NASA-DeBakey Ventricular Assist Device (VAD). NASA Ames scientists Cetin Kiris and Dochan Kwak employed NASA Shuttle main engine technology and CFD modeling capabilities, coupled with the NASA Advanced Supercomputing (NAS) Division's high-performance computing

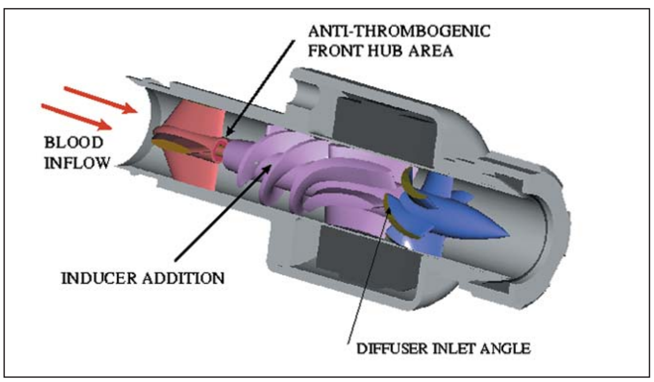

Using CFD analysis, NAS researchers found that major design modifications to the NASA-DeBakey Ventricular Assist Device were necessary. The result of these changes: overall efficiency of the device was increased by 22 percent.

(MicroMed Technology)

technology, to make several design modifications that vastly improved this miniature heart pump's performance. To date, it has been successfully implanted in 170 individuals in Europe, Asia, and the United States. In April 2002, the NASA-DeBakey VAD was named NASA's Commercial Invention of the Year for 2001.

Information Power Grid

Grids are an emerging technology that provide seamless and uniform access to geographically dispersed computational, data storage, networking, instruments, and software resources needed for solving large-scale scientific and engineering problems.

NASA's Information Power Grid (IPG) project is developing and deploying such a computing and data grid, and its goal is to use NASA's remotely located computing and data system resources to build distributed systems that can address problems too large or complex for a single site.

IPG Job Manager

The IPG Job Manager allows IPG users to reliably execute applications on remote computer systems specified by users. The Job Manager provides the following functionalities: moving files between file servers before an application begins executing (pre-staging) – a typical use of this is to move input data files to a computer system where an application will be run; moving files between file servers after an application completes executing (post-staging) – a typical use of this is to move output data files from a computer system where an application was run to a repository of project data; and reliably executes an application. The IPG Job Manager handles a variety of software and hardware failures; monitors the progress of the execution of a job; and maintains a history of the execution of an application that is available to users after the application finishes executing.

IPG Resource Broker

A computational grid such as the IPG consists of a large number and variety of computer systems. Users may not be familiar with all the available resources and their current state, so it can be difficult for users to select where to run an application. The purpose of the IPG Resource Broker is to provide suggestions for which computer system should be used to execute an application. Users provide: the number of CPUs and the amount of memory required; the executables to use for different operating systems; job execution parameters such as command line arguments, environment variables, scheduling queue, project, etc.; and the files to pre- and post-stage (optionally qualified by host name). From this input, the Resource Broker suggests where to run an application in the form of Job Manager Jobs. Each Job Manager Job contains all the information needed to execute the application, including the suggested computer system. The broker suggests systems by considering: the candidate hosts and operating systems (if any) that the user specified; the number of CPUs and amount of memory specified by the user; and the current status and load on the systems; the predicted completion times of the application on different systems.

Performance Prediction

Performance grids allow users to access a large number and variety of computer systems. This leads to the question of where users should run their applications. The NASA IPG team provides predictions of application execution and completion times to assist users in making this decision.

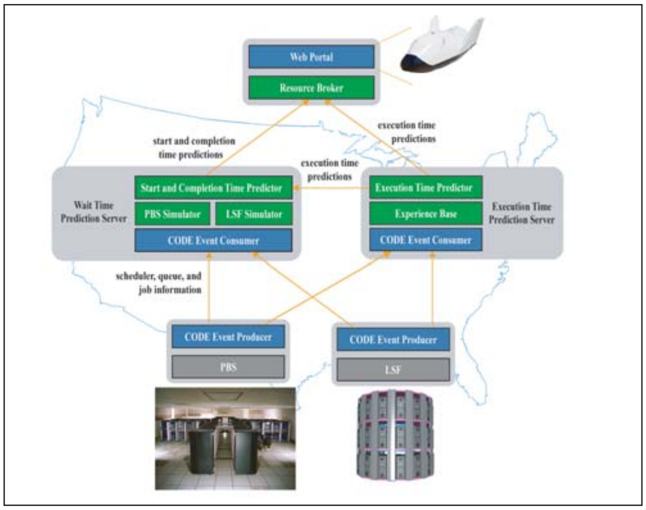

IPG's performance prediction architecture enables users to predict how an application will run on the grid, and how long the task will take. (NASA)

The team predicts execution times of applications on space-shared parallel computers using instance-based learning techniques. This tool maintains a database of experiences, called an experience base. The prediction for a query is formed using the outcomes of experiences in the database that are similar to the query. Experiences are job descriptions consisting of characteristics such as user name, number of CPUs, and executable name.

Access to most large computer systems is scheduled using queues and a scheduling algorithm that decides when to start jobs. Completion times of applications are predicted by simulating the operation of scheduling algorithms, given the jobs currently in the queues, and predictions of how long those jobs will execute.

Multi-level Parallelization and Optimization Environment

NASA researchers show three prototype tools working together to transform serial codes into multi-level parallel codes. Depicting a typical parallelization scenario, it starts with a serial cloud-modeling code and uses CAPO, a tool for assisting in parallelization, to produce an OpenMP version of the code. User inputs during parallelization are incorrect, and the resulting program does not produce the same answers as the original serial code. An automatic debugging tool is then applied, which uses information from CAPO to isolate the cause of the differences. Once the problem is fixed, the Paraver tool is used to demonstrate how to analyze performance of the parallel program. The scenario concludes with a visualization of the data computed by the program. Visit: www.nas.nasa.gov/Tools

Programmable Visualization of EOS DAAC Data

This demonstration shows programmable visualization components for the Earth Observing System (EOS) Distributed Active Archive Center (DAAC) data. The components include a central data model ("Field Model"), a metadata model ("Active Metadata"), and modules capable of interfacing to various data formats, including HDF-EOS files. The components can easily be assembled within a framework based on the Python programming language. Python includes numerous modules for web data access,

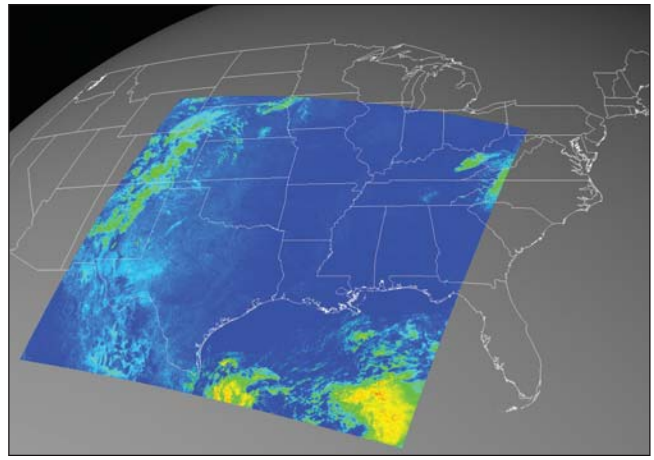

The visualization of the U.S. Gulf Coast region is based on Level One-B swath data acquired September 29, 2000, by the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument aboard the Terra satellite. Terra is the flagship in NASA's Earth Observing System. The datasets are made available by the Distributed Active Archive Center based at Goddard Space Flight Center, Greenbelt, Maryland.

offering the potential for automated access of DAAC data. The interactive visualization modules, coupled with the programmability of Python, enables users to complement the insights gained by interactive techniques with the efficiencies of programmable analysis. More information is available at: http://field-model.sourceforge.net

*-Ray: Large-Scale Data Rendering Software

Pronounced *star-ray*, *-Ray is a multi-threaded software package developed by researchers at the University of Utah and is capable of rendering large-scale scientific data sets at interactive rates. Star-ray rendering is entirely software based, taking advantage of the ccNUMA architecture of the SGI Origin series. Star-ray is able to achieve near linear scaling in performance, based on the number of pixels rendered and number of processors used. Ray tracing also provides an ideal framework for rendering large-scale data, because the rendering time scales sublinearly with scene complexity. The primary limitation on data size is main memory. With NASA's SGI Origin 3000, researchers will demonstrate the ability to interactively render datasets ranging from gigabytes to hundreds of gigabytes.

Simulation-Based Medical Planning for Vascular Surgery

The current paradigm for cardiovascular surgery planning relies exclusively on diagnostic imaging data to define the present state of the patient, empirical data to evaluate the

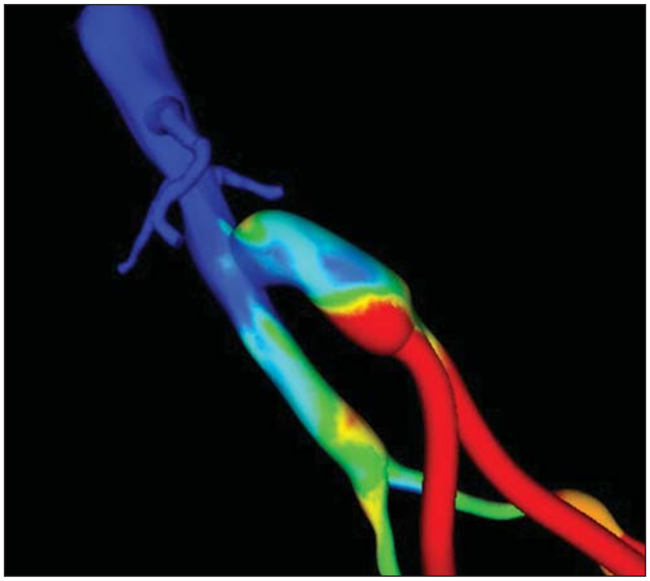

Flow simulations are performed for each preoperative plan and quantities of interest (such as wall shear stress, particle residence time) are calculated. (Stanford University)

effectiveness of prior treatments for similar patients, and the judgment of the surgeon to decide on a preferred treatment. Researchers at Stanford University and NASA are working to create a simulation-based medical planning system for cardiovascular disease that uses computational methods to evaluate alternative surgical options prior to treatment, using patient-specific models of the vascular system. The blood flow simulations enable a surgeon to see the flow features resulting from a proposed operation and to determine if they pose potential adverse effects such as increased risk of atherosclerosis and blood clot formation. Further details on this research can be found at:
www.med.stanford.edu/school/vascular

John H. Glenn Research Center

Cleveland, Ohio

As a diverse team working in partnership with government, industry, and academia to increase national wealth, safety, and security, protect the environment, and explore the universe, the John H. Glenn Research Center develops and transfers critical technologies that address national priorities through research, technology development, and systems development for safe and reliable aeronautics, aerospace, and space applications.

Coupling Engineering Applications on NASA's IPG: A CORBA-Based Approach

Researchers will demonstrate the results of an engineering parameter study that pairs NASA Langley Research Center's VULCAN (Viscous Upwind aLgorithm for Complex flow ANalysis) code with CORBA-based services and applications developed at Glenn. The demo will feature this CORBA-based environment transferring data and executing the VULCAN code on IPG resources.

A secondary benefit of this research is the ability to couple CORBA-IPG middleware to Glenn's Numerical Propulsion System Simulation (NPSS) CORBA Component Developer's Kit (CCDK). This capability enables VULCAN to function as a service to both NPSS and stand-alone IPG client applications.

Interactive Large-Screen Visualizations

NASA Glenn researchers will showcase recent activities in adapting interactive large-screen visualizations of scientific and engineering research. Demonstrations will include: examples of research in space communications and distributed supercomputing solutions using advanced propulsion models; fuel cell design and research; biotechnology; and microgravity science, all supported by NASA's Computing, Information, and Communication Technologies Program.

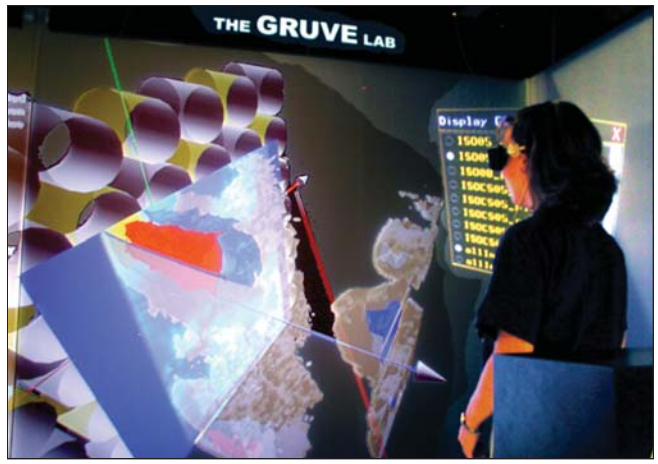

Glenn Research Center's interactive visualization theater will feature demonstrations of interactive computational research.

Goddard Space Flight Center

Greenbelt, Maryland

Goddard Space Flight Center (GSFC) seeks to expand knowledge of the Earth and its environment, the solar system, and the universe through observations from space. High-performance computing interprets observational data both by processing it into understandable forms and by simulating observed and unobserved phenomena.

Modeling the Earth's Atmosphere at the Data Assimilation Office (DAO)

The NASA finite-volume General Circulation Model (fvGCM) is the DAO's next-generation modeling system. It is based on a unique, state-of-the-art finite-volume dynamical core and community-built physical parameterizations and land surface model. A highly efficient application of the MPI-2 remote memory access message-passing paradigm has been implemented within the NASA fvGCM. This implementation improves the throughput by as much as 39 percent, compared to the MPI-1 non-blocking communication paradigm. The NASA fvGCM produces high-resolution global forecasts capable of resolving atmospheric

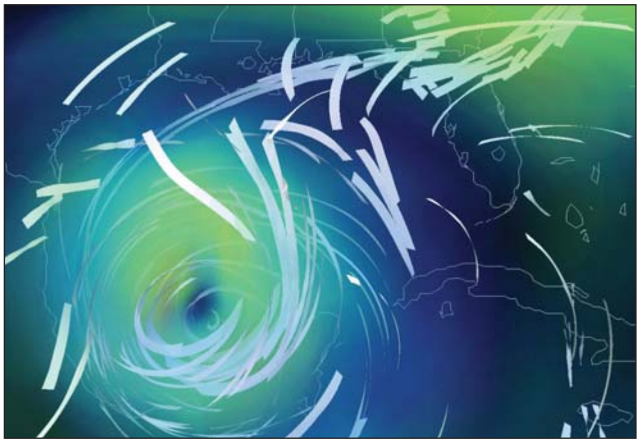

Snapshot of a three-dimensional animation of a hurricane developing in the Gulf of Mexico created from a 0.5- x 0.625-degree, 32-level run of the NASA fvGCM. (NASA)

motions from meso- to planetary-scale with a high throughput on multiple-processor, distributed memory computing platforms. For more information about the Data Assimilation Office, visit: http://dao.gsfc.nasa.gov

Ultra-High Resolution Astronomy: Phasing of Arrays of Formation Flying Spacecraft

Goddard Space Flight Center researchers are investigating arrays of multiple spacecraft, phased together and flying in formation. Systems of this type would enable ultra-high resolution science, including resolution of the stellar disks of nearby stars, as well as the planets around these stars. Systems with baselines (largest distance between spacecraft) of up to 500 meters are currently under study. GSFC scientists are researching and developing tools for systems modeling to determine the system design drivers, set science goals, and ascertain system performance. This is inherently a large-scale computing problem, both in terms of modeling and in terms of the algorithms to actively control the system. It is likely that such systems would also require autonomous on-board computing.

http://code935.gsfc.nasa.gov/cube%20Folder/OSCAR/index.html

IBEAM: Interoperability-Based Environment for Adaptive Meshes

IBEAM's technology goals are to develop a community framework for scalable astrophysical modeling simulations that promotes code interoperability. Design constraints require that the framework enable adaptive mesh refinement techniques; support interoperability of astrophysical modeling codes; be component-oriented; promote Object Oriented Design (OOD) concepts to enhance code reusability; and be coded in F90+MPI to promote performance and portability. The testbed problem is modeling gamma-

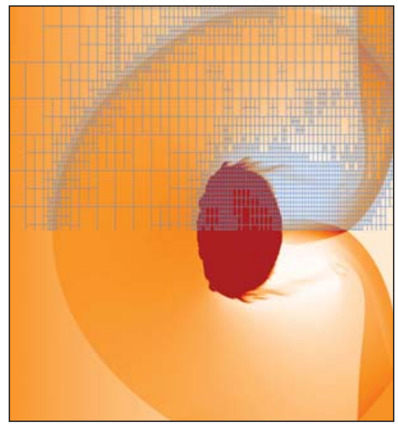

Adaptive mesh refinement simulation of a relativistic shock wave interacting with a dense gas cloud. (NASA)

Chombo Framework for Block-Structured AMR Applications

Researchers are developing a software framework for implementing block-structured adaptive mesh refinement (AMR) algorithms to solve targeted problems in star formation, microgravity research, and space physics. The Chombo framework is a basis for this effort. Chombo is a set of object-oriented C++ libraries for implementing AMR applications that use MPI to support distributed memory computing.

Algorithm components in development include finite-volume techniques for representing irregular geometries and multi-fluid interfaces, along with techniques for particles in incompressible fluids.

The ChomboVis tool provides an interpretive environment for manipulating and visualizing AMR data in the Chombo framework. ChomboVis is built using the Python scripting language and the Visualization Tool Kit, enhancing portability to a variety of architectures. Both Chombo and ChomboVis use HDF5 (Hierarchal Data Format) for input and output. For more details, visit: http://seesar.lbl.gov/ANAG/NASA

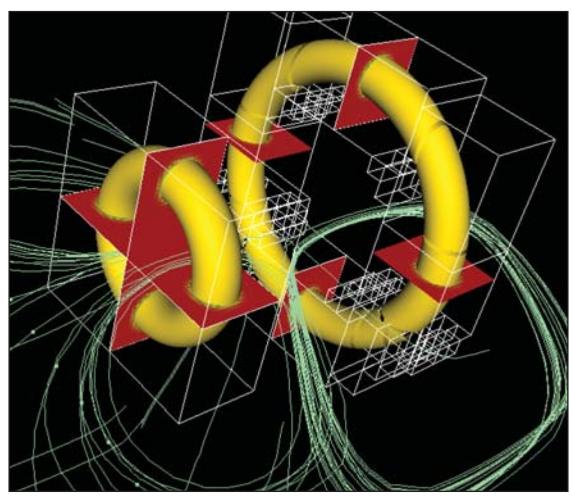

Adaptive mesh refinement (AMR) simulation of two co-rotating vortex rings (indicated by the yellow rings). Green lines show streamlines of the flow, while white boxes show where finer resolution is used to better capture details of the flow. Coherent vortices like those shown here are the principal means for mixing fluids in microgravity environments.

ray burst fireballs and afterglows. Project participants include NASA GSFC, University of Illinois, University of Chicago, SUNY Stony Brook, and the Universities Space Research Association. Details of this research can be found at: www.ibeam.org

The Parallel Virtual File System Version 2

PVFS is a parallel file system for Linux clusters. It harnesses commodity storage and network technology to provide concurrent access to data that is distributed across a (possibly large) collection of servers. PVFS serves as both a test-bed for parallel I/O research and as a freely available, production-level parallel file system for use in the cluster community. Development and support is carried out with collaboration between Clemson University, GSFC, and Argonne National Laboratory. For more information, visit: www.parl.clemson.edu/pvfs

Coven Software Development Framework

Coven is a framework for component-based collaborative software development. Its goals are to provide extensibility and usability, and to produce high-performance code. Extensibility is achieved through pluggable modules representing individual algorithms. Modules are connected to form a dataflow graph, created with assistance from a graphical user interface. Environments using this framework can build additional graphical components, which tailor to the specific needs of that problem domain (such as dataset visualizers, performance analyzers, debuggers, database query tools). Coven has been used to build environments for satellite remote sensing, n-body problems, molecular dynamics problems, and CFD fluid flow and heat transfer problems. For more information, visit: www.parl.clemson.edu/Coven

Beowulf/Mini-Grid System Software

A mini-grid is a campus-wide version of a computational grid, featuring a private switched network interconnecting grid resources where performance is not limited by Internet connections. All resources on the Clemson grid are Beowulf/Linux clusters. The Scyld Beowulf Linux OS has been modified to support the mini-grid with the addition of Balloc, the Beowulf node Allocator, which allows large jobs to transparently "borrow" nodes from adjoining clusters. Traffic for jobs on borrowed nodes runs through a private network, which means performance can be predicted. A single job can use all processors in the grid. The software project is funded by Goddard. Details are on the web at: www.parl.clemson.edu/minigrid

Cost-Effective Advanced Computing Technologies

Researchers at NASA are working on the use of reconfigurable logic to provide ultra-high-performance computing in both space-based and ground-based systems. Techniques include implementing applications directly in reconfigurable hardware, implementing specialized assist logic in reconfigurable hardware, and using reconfigurable logic blocks inside of an otherwise standard microprocessor. Information on this NASA/Clemson University collaboration is available at: www.parl.clemson.edu

Sci-Interactives: Science Outreach from Truth-N-Beauty Software and NASA

The advent of broadband Internet access together with high-speed processors for commercial PCs, has opened significant new avenues for communicating science to the general public. Truth-N-Beauty Software, a science 'e'- outreach and education company, has partnered with NASA's Computational Technologies (CT) Project to develop digital tools for teaching the public about the latest results of NASA Earth and space science research.

Sci-Interactives are web-based applications that let the science speak for itself. These tools give people access to the real equations and data behind NASA science with intuitive, graphically sophisticated tools and games.

Current projects include story development and Sci-Interactives for *Discover* Magazine's website. Tools explore the collapse of rotating interstellar clouds and the creation of aurora in the Earth's magnetosphere. Details of NASA's latest Earth and space science research can be found at www.truth-n-beauty.com

Simulation screen for a collapse and rotation Sci-Interactive.

Jet Propulsion Laboratory (JPL)

Pasadena, California

The Jet Propulsion Laboratory's mission is to enable the nation to explore space for the benefit of humanity. The lab's mission is: to explore our own and neighboring planetary systems; to search for life outside the Earth's confine; to further our understanding of the origins and evolution of the Universe and the laws that govern it; to make critical measurements to understand our home planet and help protect its environment; to apply JPL's unique skills to solve problems of national security and national significance; and to inspire the next generation of explorers.

Shuttle Radar Topography Mission(SRTM): Science Processing on a Supercomputer

NASA researchers have developed software to process raw data acquired by the Shuttle Radar Topography Mission (SRTM), as part of an effort to generate a global digital elevation model database. The large amount of data required to form the mosaic was processed on *Alhena*, the SGI Origin 2000 supercomputer at JPL, using 32 processors in parallel. Digital elevation models are used for scientific applications in geophysics, geology and hydrology, as well as practical applications, such as highly detailed 3-D maps for civilian and military user communities. Information about this project, co-sponsored by NASA and the National Imagery and Mapping Agency, can be viewed at: http://photojournal.jpl.nasa.gov.

Dusty Rings: Signposts of Recent Planet Formation

Every planetary system forms in a thin disk of gas and dust around a young star. Web-based animations show the evolution of bright rings of dust in a disk where planets form. Planets grow from collisions and mergers of smaller bodies, called planetesimals, embedded in a disk of gas and dust in orbit around the star. Each visualization starts with the same initial conditions, roughly 10 billion billion planetesimals 1-100 meters across, embedded in a disk that extends from the orbit of Neptune to the outer edge of the Kuiper belt. The colors represent the intensity of the dust grains from low (dull reds) to high (blue/white). Individual ani-

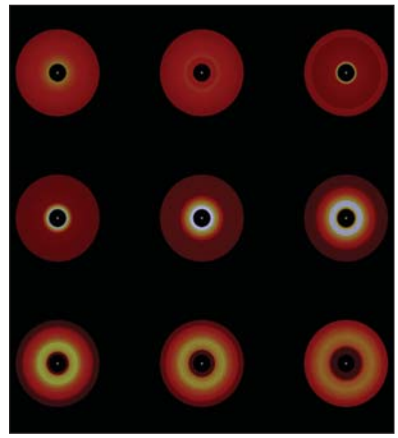

Estimating the Circulation and Climate of the Ocean: The Circulation Pathway of Subtropical-Tropical Exchange

JPL researchers are investigating mechanisms of subtropical-tropical water mass exchange in the Pacific Ocean, using an animation of a simulated passive tracer and its adjoint tracer. The forward passive tracer and adjoint passive tracer can be identified as describing where the tracer-tagged water mass goes to and where it comes from, respectively. The exchange plays a central role in the decadal change of *El Niño*, and is examined by explicitly tracing the water of the eastern Equatorial Pacific. In contrast to previous theoretical expectations, the exchange is found to be an "open-circuit," in which the origin is distinct from the destination. Temporal variability in ocean circulation is found to be fundamental in determining the pathway of the water mass, "short-circuiting" the otherwise circuitous route. Details of this research are available on the Internet at: http://ecco.jpl.nasa.gov/external

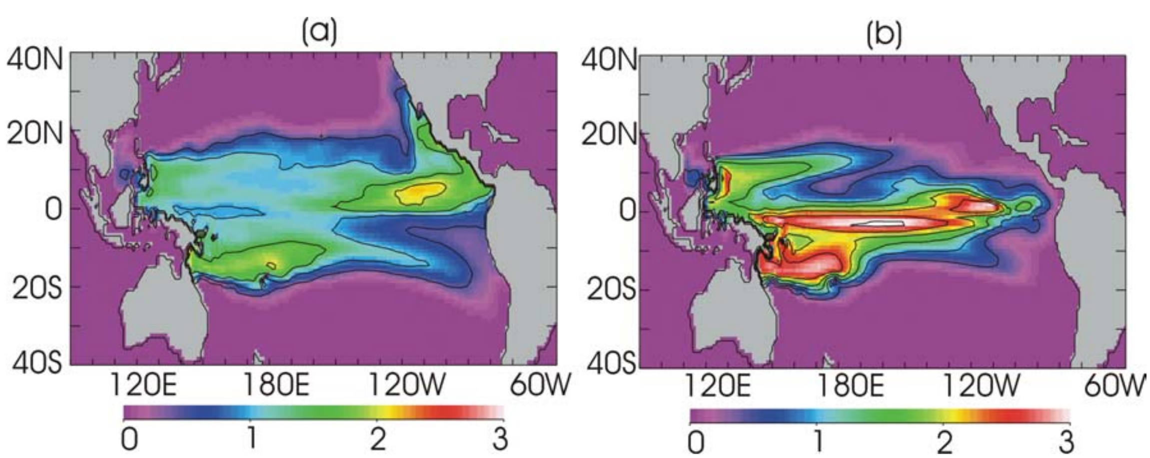

The image at left illustrates the average distribution of the water mass (depth integrated content) that reaches the eastern Equatorial Pacific five years prior to its terminal instant. The right panel shows the same data, except when the time-mean circulation is assumed. The difference between the two illustrates the significance of the circulation's temporal variability.

mations depict a passage of time of 2 billion years. In the "weaker bodies" animation, the bodies fragment easier than the stronger bodies, and thus produce more dust and brighter rings when they collide. The figure shown on page 8A depicts nine snapshots of the disk in one calculation. The bright point at the center is the star. The inner edge of the disk is at the orbit of Neptune. The outer edge is at the outer edge of the Kuiper Belt in our solar system. The top row of frames corresponds to a time, t = 0, 3, and 10 million years (Myrs); the middle row is t = 30, 100 and 300 Myrs; the bottom row is t = 600 Myrs, 1 billion years (Gyrs) and 2 Gyrs. Further details can be found at: http://cfa-www.harvard.edu/~kenyon/pf/sp/index.html

Using Hollywood for Science: Animation of a Proposed NASA Mission using Maya 4.0

The Satellite Test of the Equivalence Principle (STEP) was a proposed NASA small explorer mission which would test the principle that all masses, regardless of size or composition, will experience the same acceleration in a given

Continued on page 10A

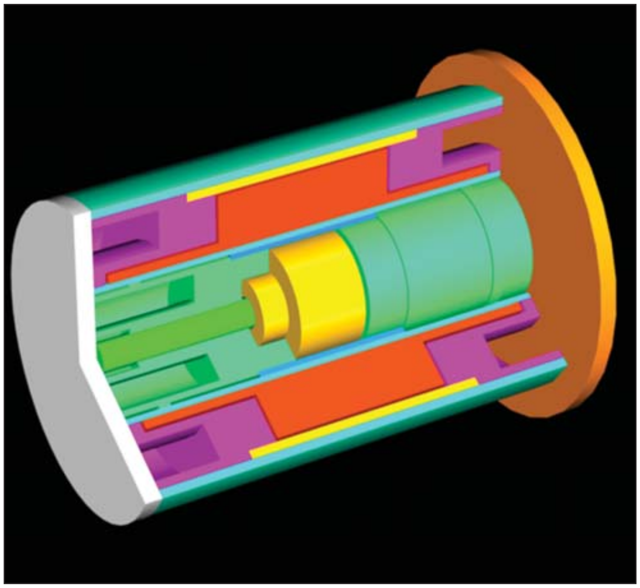

Using Hollywood for Science: Animation of a proposed NASA mission using Alias Wavefront's Maya 4.0 software.

3-D Adaptive Grid: MHD Simulations of the Global Heliosphere

The visualizations show the interaction of the interstellar medium and the solar wind. The solar system is moving through the local interstellar medium. This complex system involves three different discontinuities: the termination shock, the heliopause, and the bow shock. In the visualization, the interstellar wind flows from the left and impinges on the solar wind. The sun is located at the center.

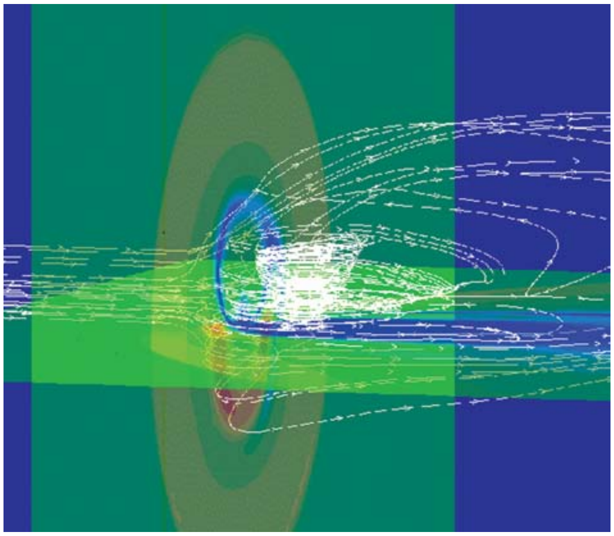

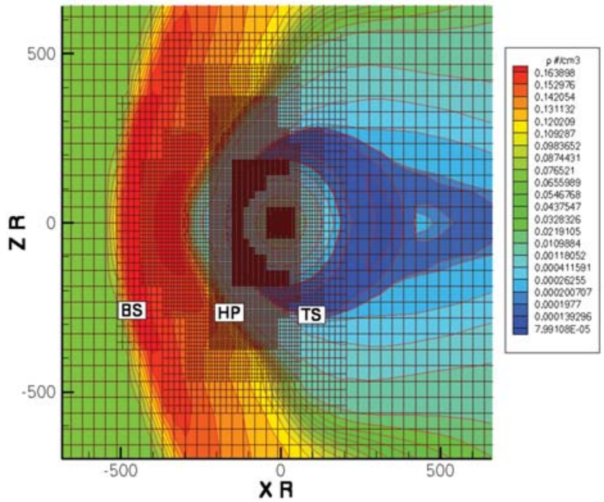

Above: *This graph shows the contours of the density. In this figure the interstellar medium flows from the left, with the Sun at the center of the figure. Black denotes different refined areas of the mesh. Also indicated: the location of the Bow shock (BS), Heliopause (HP), and the Termination shock (TS).*

Left: *Contours of the magnetic field are shown with blue representing the minimum intensity of the magnetic field and red depicting maximum intensity. The white lines indicate magnetic field lines. At the center, the solar magnetic field (the Parker spiral), can be seen with the Sun at the center. The interstellar medium flows from the left.*

The code uses BATS-R-US, developed at the University of Michigan. It is a 3-D magnetohydrodynamics adaptive grid code. It uses an adaptive Cartesian grid made of a rectangular block of cells. The solar wind and the interstellar wind are modeled, including both magnetic fields and the plasma in a self-consistent way. The simulation used 4.5 million cells and 11 levels of refinement. The minimum size of the cells is 0.5 Astronomical Units (AU), and the maximum size is 36 AU.

Continued from page 9A

gravitational field. This principle is a fundamental assumption in Einstein's General Theory of Relativity. The movie contains a simulation of the relative motion of two test masses in the experiment in orbit about the Earth when the Equivalence Principle is violated, followed by an animation of the assembly of the various parts of the experiment. The modeling, animation, and rendering of the movie were performed using Alias Wavefront's Maya Complete software package, and took approximately three hours on eight processors of JPL's SGI Origin 2000 supercomputer. Further details on this research project can be found at:

http://sc.jpl.nasa.gov/success_stories/slides/slide44.html

Common Component Architecture: A Demonstration of the CCA Forum Technologies

The Common Component Architecture (CCA) Forum is a group of researchers from national labs and academic institutions committed to defining a standard component architecture for high-performance computing.

In the last year, researchers began testing the Common Component Architecture (CCA), designed for large, collaborative, high-performance computing applications. The testing is designed specifically for its suitability to NASA applications, such as climate simulations. Further details can be found at: http://pat.jpl.nasa.gov/public/projects

NASA Participation: Technical Papers, Tutorials, and Panel Discussions

Several NASA scientists (both contractor and civil service personnel) are participating on SC2002 committees, panels, tutorials, and other conference events.

Paper Presentation

Collaborative Simulation Grid: Multiscale Quantum-Mechanical/Classical Atomistic Simulations on Distributed PC Clusters in the USA and Japan

Hideaki Kikuchi (Louisiana State University), Rajiv K. Kalia, Aiichiro Nakano, Priya Vashista (University of Southern California), Hiroshi Iyetomi (Niigata University), Shuji Ogata, Takahisa Kouno (Yamaguchi University), Fuyuki Shimojo (Hiroshima University), Kenji Tsuruta (Okayama University), and Subhash Saini (NASA Ames).

A multidisciplinary, collaborative simulation has been performed on a grid of geographically distributed PC clusters. The multiscale simulation approach seamlessly combines atomistic simulation based on the molecular dynamics (MD) method and quantum mechanical (QM) calculation based on the density functional theory, so that accurate but less scalable computations are performed only where they are needed. The multiscale MD/QM simulation code has been grid-enabled using a modular, additive hybridization scheme, multiple QM clustering, and computation-communication overlapping. The 'gridified' MD-QM simula-

Langley Research Center

Hampton, Virginia

In alliance with industry, other agencies, academia, and the atmospheric research community, Langley Research Center (LaRC) undertakes innovative, high-payoff aerospace and scientific activities beyond the risk limit or capability of commercial enterprises and delivers validated technology, scientific knowledge, and understanding of the Earth's atmosphere. The center's success is measured by the extent to which its research results improve the quality of life.

Biological Metaphors for Feature Extraction in Large Datasets

This demonstration describes a new visualization metaphor for quickly extracting and visualizing features in large datasets. The data swarming metaphor is a biologically based collection of agents that move through space interacting with local data. Although each agent follows a few simple rules, the collection of agents can find complex patterns in the dataset. In particular, this demonstration portrays a general agent architecture that can be used to create unique visualization tools as well as support common visualization techniques. The data swarming paradigm is applied to computational fluid dynamics visualization (below), and is implemented as a virtual environment that runs on immersive displays.

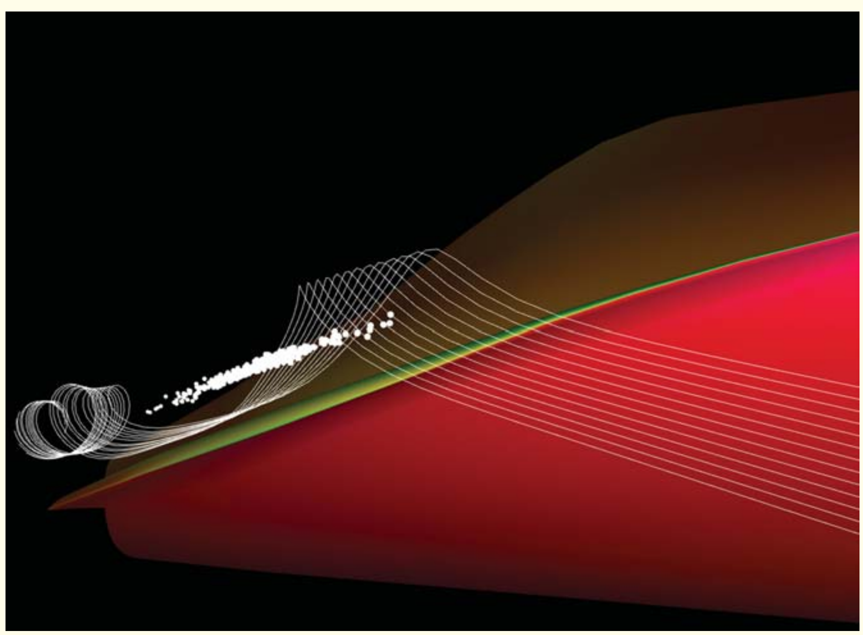

tion code has been used to study environmental effects of water molecules on fractures in silicon. A preliminary run of the code has achieved parallel efficiency of 94 percent on 25 PCs distributed over three PC clusters in the United States and Japan. A larger test involving 154 processors on five distributed PC clusters is in progress.

Tutorial S11

November 18, 8:30 a.m. – 12:00 p.m.

Requirements Analysis for Scientific Applications

Scientific applications are increasingly complex and collaborative. In order to coordinate, plan, and validate scientific software, it is useful to gather requirements before development begins. In this tutorial, developers of the Earth System Modeling Framework explore how requirements analysis can improve development efficiency and reduce misunderstandings and defects. They will look at how requirements collection can be implemented for distributed collaborations; review sample requirements documents and templates; and discuss how the requirements process impacts other aspects of software development, such as design and verification. Throughout this process they will emphasize sensitivity to the unique aspects of the research-oriented HPC computing environment. More information about the Earth System Modeling Framework is at: http://www.esmf.ucar.edu

Tutorial S3

November 16, 8:30 a.m. – 5:00 p.m.

Component Software for High-Performance Computing: Using the Common Component Architecture

Rob Armstrong, Jaideep Ray (Sandia National Laboratories), David Bernholdt, Wael Elwasif, Jim Kohl (Oak Ridge National Laboratory), Lori Freitag, Lois McInnes, Boyana Norris (Argonne National Laboratory), Dan Katz (NASA–Jet Propulsion Laboratory), Gary Kumfert (Lawrence Livermore National Laboratory), Craig Rasmussen (Los Alamos National Laboratory)

This tutorial will introduce the Common Component Architecture (CCA) at both conceptual and practical levels. Components are software objects that encapsulate useful functionality and interact with other components only through well-defined interfaces. Component-based approaches to software development, including the CCA, provide a means to manage the complexity of large-scale software applications and facilitate the reuse and interoperability of code.

The component concept is widely used in the commercial and business world (e.g., CORBA, COM, and Enterprise JavaBeans) and is popular in visualization as well (e.g., Data Explorer), but none of these environments is well suited to use in high-performance computing environments. The CCA was designed specifically with the needs of high-performance scientific computing in mind, including the existence of large bodies of existing code that cannot easily be rewritten. The tutorial will cover the concepts of components and the CCA in particular; the tools provided by the CCA environment – including the Babel system for language interoperability; the creation of CCA-compatible components; and their use in scientific applications. A combination of traditional presentation, and live demonstration will be used during the tutorial. After the session, participants will be able to download the tools and software to reproduce the live examples in their local environment. For more information, see:
www.cca-forum.org/tutorials/sc02.html.

Panels

November 22, 8:30 – 10:00 a.m.

Are Designer Supercomputers an Endangered Species?

Moderator: *Aruna Ramanan (IBM)*

Panelists: *Thomas Sterling (Center for Advanced Computing Research, California Institute of Technology and Jet Propulsion Laboratory); Gita Alaghband (University of Colorado, Denver); Jamshed Mirza (IBM Corporation); Tadaski Watanabe (NEC Corporation); and Candace Culhane (NSA)*

At the dawn of the high-performance computing era, all supercomputers were “designer machines.” Every aspect of these machines was designed specifically to lend the highest performance level achievable, using the state of the art technology available at the time. While the attention to design still remained, the focus changed from exploiting fine-grained parallelism using vector processors to coarse-grained parallelism using “massive parallel processors” or MPPs. The emergence of MPPs changed the focus from the processor to the interprocessor communication network. The systems were still designer systems. Then came the shift to specially designed clusters.

More recently, “Components Off the Shelf” or COTS clusters, have emerged in a big way. Strides in the cyber communication infrastructure, together with the narrowing gap between the computing world and the networking world has opened up the possibility of grid computing and peer-to-peer computing. The SETI@home project has demonstrated the power of utilizing distributed resources. These developments have led many to believe less and less in designer supercomputers and more and more in constructing high-performance systems using widely available components used in mainstream computing. But the emergence of the Japanese Earth Simulator as the top performer has stirred up the high-performance computing community. This panel will explore the strategies needed to transform terabytes of information into insights, and examine the challenges that need to be addressed by the two different approaches to supercomputing.

SC2002

BOOTH MAP

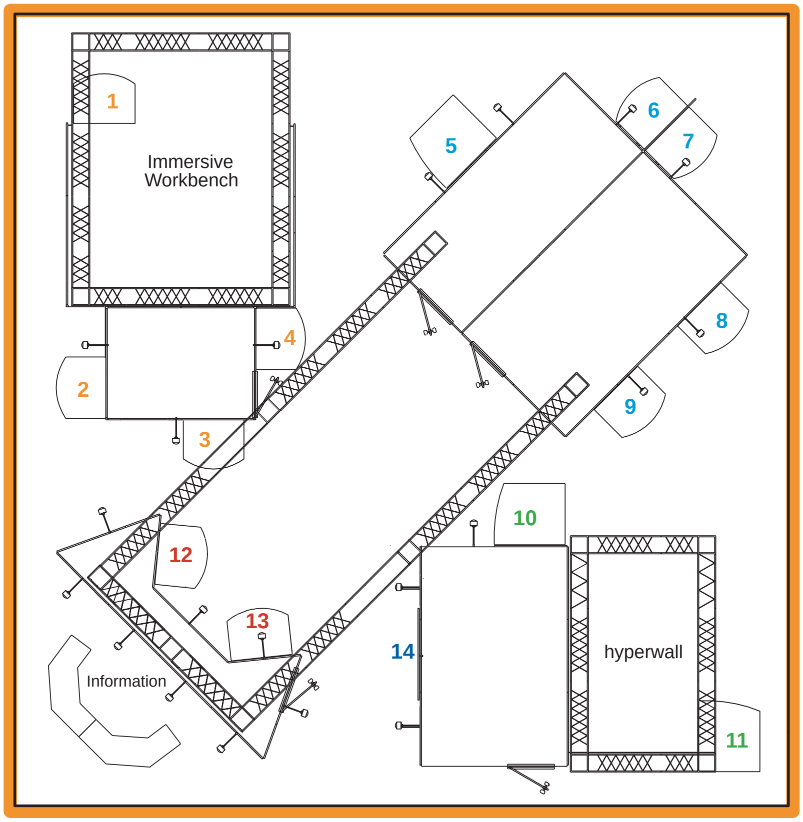

1 Immersive Workbench
- Biological Metaphors for Feature Extraction in Large Datasets
- IBEAM: Interoperability-Based Environment for Adaptive Meshes
- Interactive Large Screen Visualizations

2
- Coupling Engineering Applications on the Information Power Grid: A CORBA-Based Approach
- 128-Node Information Power Grid Cluster

3
- Simulation-Based Medical Planning for Vascular Surgery

4
- Common Component Architecture (CCA): A Demonstration of the CCA Forum Technologies
- Dusty Rings: Signposts of Recent Planet Formation
- Estimating the Circulation and Climate of the Ocean: The Circulation Pathway of Subtropical-Tropical Exchange
- Shuttle Radar Topography Mission (SRTM): Science Processing on a Supercomputer
- 3D Adaptive Grid: MHD Simulations of the Global Heliosphere
- Using Hollywood for Science: Animation of a Proposed NASA Mission Using Maya 4.0

5
- AeroDB: Automated CFD Parameter Studies on Distributed Parallel Computers
- Information Power Grid (IPG)
- IPG Job Manager
- IPG Resource Broker
- Performance Prediction
- Tool Agent Framework-Java (TAF-J) - Connecting Aerospace Analysis Tools with the Information Power Grid

6
- Multi-Level Parallelization and Optimization Environment

7
- Atomistic Simulations on USA and Japan Grid
- Cart3D: A Package for Automated Grid Generation and Aerodynamic Database Creation

8 & 9
- Beowulf/Mini-Grid System Software
- Coven Parallel Programming Framework
- The Parallel Virtual File System Version 2
- Reconfigurable Computing Technology

10
- Programmable Visualization of EOS DAAC Data
- Real-Time Ray Tracing with the University of Utah

11
- The hyperwall: A Multiple Flat-Panel Display Wall for High-Dimensional Visualization

12
- Modeling the Earth's Atmosphere at the Data Assimilation Office
- Sci-Interactives: Innovative Science Outreach from Truth-N-Beauty Software and NASA CT
- Ultra-High-Resolution Astronomy: Phasing of Arrays of Formation Flying Spacecraft

13
- Chombo Framework for Block-Structured AMR Applications
- IBEAM: Interoperability-Based Environment for Adaptive Meshes

14 Video and Poster Wall
- SC2002 Conference Opening Video
- Capability Computing at the NASA Ames Research Center: Solving the Toughest Computational Problems
- Cart3D: 2002 NASA Software of the Year
- Computational Science at Goddard Space Flight Center
- Computing, Information, and Communications Technology Program (CICT)
- Computing, Networking, and Information Systems Project (CNIS)
- Consolidated Supercomputing Management Office (CoSMO)
- HiMAP: High-Fidelity Multidisciplinary Analysis Process
- Information Power Grid (IPG) Accomplishments - 2002
- MU-SPIN: Extending Computing Opportunities to Minority Institutions
- NASA-DeBakey Ventricular Assist Device (VAD) - 2001 NASA Commercial Invention of the Year
- NASA Research Highlights Video
- NASA/SGI Develops 1,024-Processor Coherent Shared-Memory Origin 3000
- Numerical Propulsion System Simulation (NPSS)
- Supercomputing at Goddard Space Flight Center

The Vision Behind the 'hyperwall'

NAS researchers have developed an innovative visualization system that allows scientists to view complex datasets and multiple parameters.

The flow field around an aircraft as it soars through the skies. A dynamic model of cloud formation. A molecule undergoing bond-breaking in quantum mechanical fields. All of these complex simulations are typical examples of how today's supercomputers are generating more data than ever before. And as the quantity of available data increases, advances in computational science have improved the quality of that data. Whereas computing the flow around an aircraft at a given speed and altitude was once considered a large and sophisticated process, today's supercomputers can readily compute the flow around a whole family of vehicle designs, with a wide range of different speeds, altitudes, flight profiles, and other variables.

A team of NAS scientists has developed a new tool to help researchers display, analyze, and study these high-dimensional datasets in meaningful ways – the hyperwall. Inspired by the sight of endless panels of images pinned to researchers' walls, desks, and corridors, the hyperwall is a seven-by-seven cluster of flat panel screens, each driven by its own dual-processor computer with high-end graphics card. Unlike powerwalls or caves – large projection systems that have to coordinate separate machines to tile images, the hyperwall allows each computer to process and display its own data. This essential difference enables researchers to harness that computer power and provides the flexibility to display data, while allowing their own human perceptual apparatus – the eye-brain system, to perform the functions it does best – feature detection, pattern recognition, and sorting.

"The hyperwall provides a problem-solving environment where we can use different tools, different viewpoints, even different graphics parameters, to display the same data or relationships within a dataset," explains the system's developer, NAS Senior Scientist Chris Henze. Together with NAS scientists Timothy Sandstrom, Creon Levit, and David Ellsworth, Henze received a grant from the Information Sciences and Technology Directorate at NASA Ames in April 2002 to put the system together and demonstrate the basic proof-of-principle. Today, the hyperwall's power and versatility are already attracting attention from research teams both within and outside the NASA Ames community.

The hyperwall team recently demonstrated that power by showing a visiting group of researchers from NASA's Goddard Space Flight Center how the hyperwall can be used to display Goddard's models of cloud data. "The anatomy of clouds is incredibly complex," explains Henze. "Goddard's cloud models involve a wide range of hydrometeor classes (different forms of water and ice), which intermingle within a cloud. If you try to show them in a single image, it just looks a mess. On the hyperwall, we were able to show the researchers separate classes on different screens along one column of the wall, with different time steps running from left to right, so they could see the evolution of these different fields."

Although the hyperwall can display a single image distributed across all of its screens, the team has focused on developing software to enhance the system's ability to display sets of related images. Its strength, they believe, lies in its ability to take one simulation and look at many different quantities, cross-sections, magnifications, and time steps, while keeping the whole image, or images, visible in subarrays, for example on three-by-three screens in the corner. "It's like a visualization spreadsheet. The array of displays is a giant spreadsheet with the user defining which operations happen in each cell, or row, or column," says Henze.

Why Seven-by-Seven?

"And finally, what about the magical number seven? What about the seven wonders of the world, the seven seas, the seven deadly sins, the seven daughters of Atlas in the Pleiades, the seven ages of man, the seven levels of hell, the seven primary colors, the seven notes of the musical scale, and the seven days of the week? What about the seven-point rating scale, the seven categories for absolute judgment, the seven objects in the span of attention, and the seven digits in the span of immediate memory? For the present I propose to withhold judgment. Perhaps there is something deep and profound behind all these sevens, something just calling out for us to discover it. But I suspect that it is only a pernicious, Pythagorean coincidence."

George Miller

The Magical Number Seven, Plus or Minus Two; Some limits on our Capacity for Processing Information

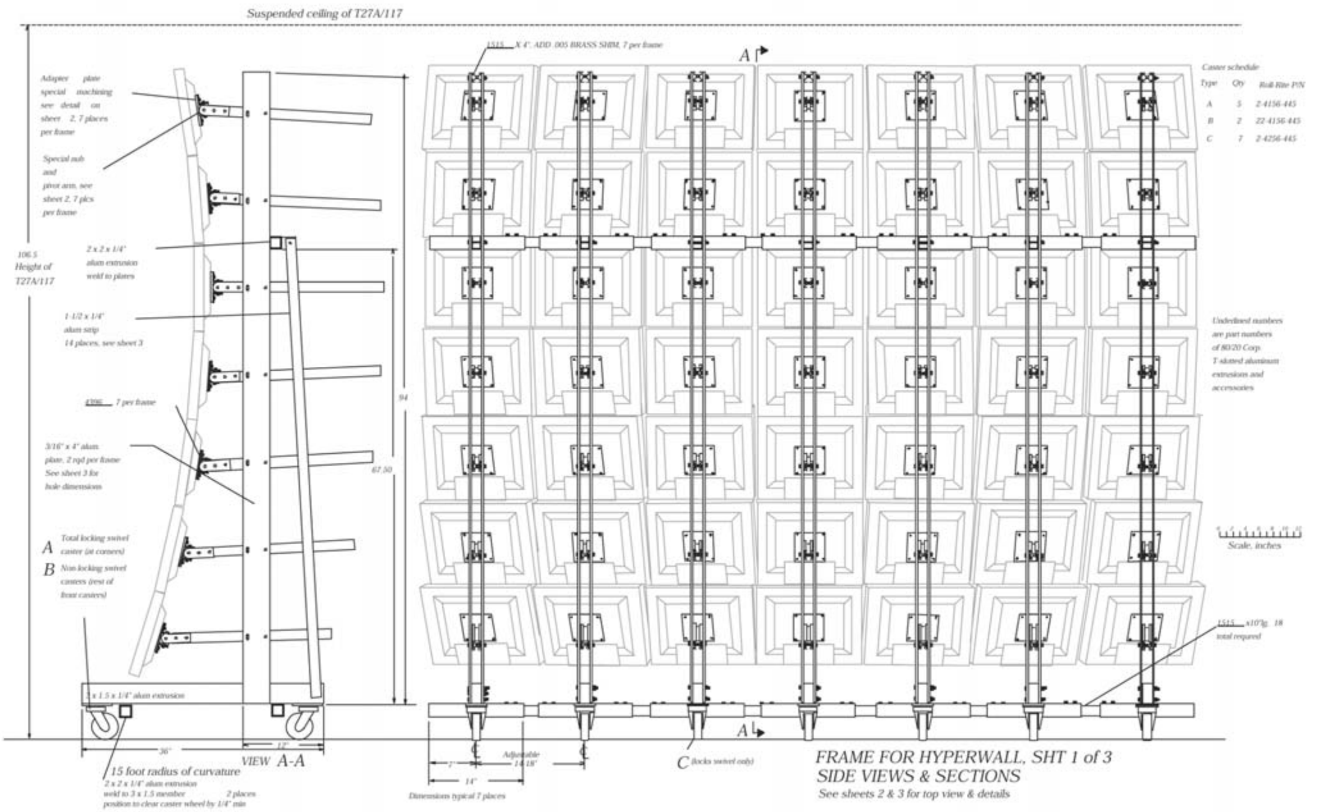

The NASA Ames hyperwall consists of 49 flat panel displays driven by an equal number of rack-mounted dual-CPU nodes. The tabular layout of the hyperwall supports many existing "multiview" visualization models, including spreadsheet-style approaches, multidimensional/multivariate techniques, and brushing/linking. (Art Joly)

To perform these endless combinations of tasks, the hyperwall uses a standard Linux operating system on the master node and on each CPU, or slave node. The master controls each screen by broadcasting script commands to copies of the programs running on each of the slave nodes. Using the team's 'growler' framework (see *Gridpoints*, Fall 2000, page 11), a system designed for distributed programming, the user is able to determine the location of processes and images with the master computer or the hyperwall serving as the graphical user interface. In addition, with a 100-gigabyte disk on each of the 49 slave nodes and even more on the master, the system has approximately five terabytes of aggregate disk space and more than 50 gigabytes of RAM.

As the project moves beyond the proof-of-concept stage, members of the hyperwall team are developing new ideas for mapping, manipulating, and managing data, including an enhancement that will allow different cooperating simulations to communicate, share, and display their composite information. But the team is also sensitive to their need to talk to researchers and to understand the application requirements.

"Visualization researchers are notorious for telling scientists the best way to display their data, when all the scientists want is a 2-D plot," says Creon Levit. "Our approach is to understand their demands exactly, ask questions, and see how we fit in. The hyperwall team has a wide range of experience and expertise that not only covers graphics, programming, and scientific visualization, but also includes research fields such as biochemistry, nanotechnology, fluid dynamics, and quantum mechanics."

This multidisciplined approach was instrumental in determining some of the human factors behind the seven-by-seven array. With an eye to George Miller's 1956 paper, *The Magical Number Seven, Plus or Minus Two: Some Limits on our Capacity for Processing Information*, as well as budget and space constraints, the team discovered that, at a comfortable distance, the 49 screens fit naturally in the overlap field of vision. "We were originally going for eight-by-eight," explains Levit. "Computer scientists naturally gravitate toward powers of two, but at a natural distance, this seven-by-seven array of monitors has a resolution of just over 64 million pixels – that's about the same resolution as your eye. So, if you have perfect vision in a dark room, theoretically you could almost tell the difference between single pixels."

Another human factor considered by the team is the issue of a seamless display. Powerwalls typically use several projectors to display a single image, and try to minimize the seams by lining up pixels and balancing colors. However, such systems are often difficult to calibrate and, although the borders are less obvious than the one-inch or so bezel around the hyper-

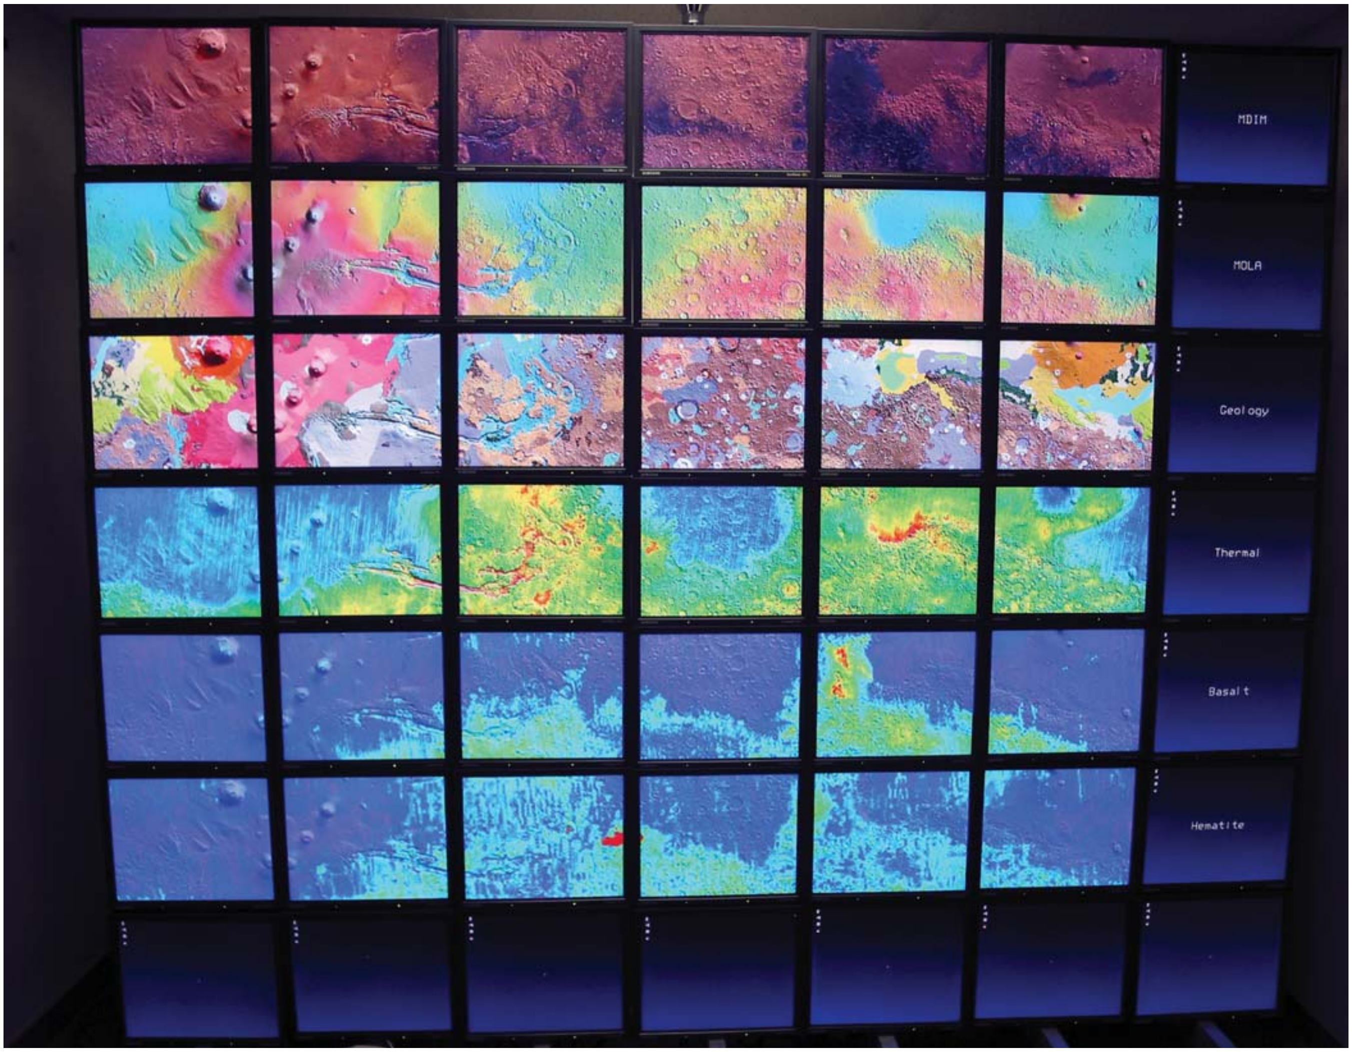

The hyperwall emphasizes displaying multiple independent, but related, images, and providing useful means for composing and controlling these image sets. In place of elaborate software or hardware crossbar switches, scientists rely on the human visual system for integration, synthesis, and pattern discrimination in complex and high-dimensional data spaces. Recently, the hyperwall was used to display Mars surface terrain images. Each row shows the same 48-degree swath along the equator from 180 degrees West to 180 degrees East. The rows show various data maps of that swath, and from top to bottom display: 1) the visible surface, 2) surface elevations, 3) geologic units, 4) thermal inertia (analogous to granularity), 5) basalt distribution, and 6) hematite distribution. This visual layout facilitates rapid intercomparison of features in different data maps. (Dataset courtesy MOLA Science Team, D. H. Scott et al., M. Mellon et al., J. L. Bandfield et al., P. R. Christensen et al.; Glenn Deardorff, NASA Ames Research Center)

wall's screens, the team found that faint boundaries were often more bothersome. "Compare the difference between a poorly constructed picture window with seams and ripples, versus a multi-paned window," says Levit. "People don't mistakenly think that the outside world is really broken up when they look through a multi-paned window."

Each screen is mounted on a custom-designed rack, and can be adjusted individually or by column or row to accommodate different viewing requirements and distances. The rack's design also allows users to compensate for the monitors' directional field of view, or make other viewing arrangements, such as gently curved display areas. And because the entire system is based on commodity technology rather than highly specialized and expensive projection equipment, the hyperwall team is confident the system will be easy to scale in the future and will decrease in cost.

In the meantime, the team continues to focus on providing compelling problem-solving environments for various research teams and pursuing software development based on individual application requirements. "We have somewhat of a symbiotic relationship with scientists," explains Sandstrom. "They need our tools, and we need their data. Their very large and complex datasets drive the capabilities of our software, so we have this leapfrog approach. They challenge us with bigger and more complex data, and we respond with better techniques. The closer computer scientists can work with the research scientists, the better results we can get for everyone."

—Julie Jervis

New NAS Parallel Benchmarks Release Addresses Grid Computing

By Rob Van Der Wijngaart

One of the NAS Division's widely recognized research products is the NAS Parallel Benchmarks (NPB). The division recently released a major new version of the benchmarks, which incorporates several important aspects of computing not covered by the previous release (see page 19). Included are benchmarks for new languages, disk access (input/output, I/O), and grid computing.

First conceived in 1991, NPB is a collection of standard tasks that are representative of classes of computational applications important to NASA and to its research and industry partners. An important NPB goal was to provide a level playing field for the myriad of high-performance computer architectures and parallelization paradigms and tools available at that time. Organizations like NASA needed to evaluate the performance of such systems without giving undue advantage to any particular vendor or tool developer, and to provide sensible targets for all. Therefore, NPB was originally formulated as paper-and-pencil problems, with very few restrictions on how they ought to be implemented and executed. The computing community embraced this concept, and submitted many results for evaluation by NAS.

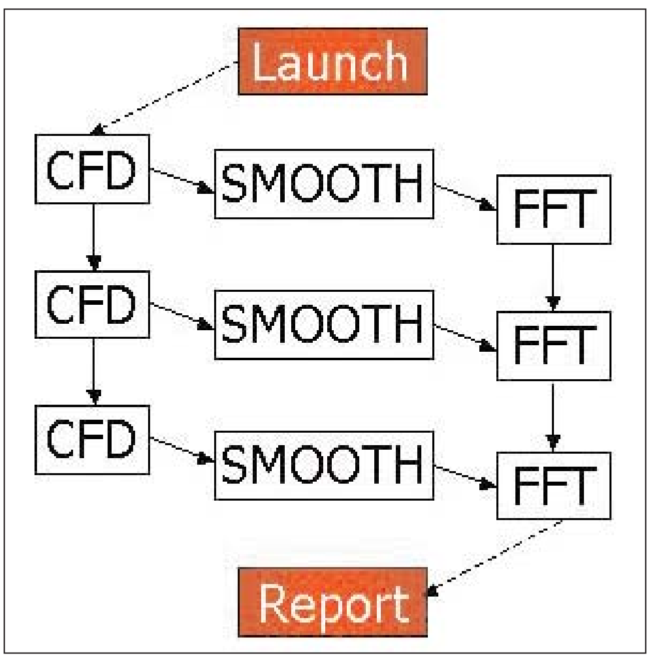

Figure 1: *Schematic of Visualization Pipe, one of four grid benchmark problems. Computation (in this case computational fluid dynamics or CFD), post-processing (Smooth), and visualization (FFT) are carried out cyclically. (Van Der Wijngaart)*

In the mid-1990s, computer architectures began to converge and standards for parallel computing emerged, so NAS researchers decided to develop an implementation of NPB that could be executed on high-performance computers worldwide. This implementation, predominantly in Fortran and using the first widely accepted parallelization standard, Message Passing Interface (MPI), was released in 1996 under the name NPB2. Several versions followed, and the final version of all eight benchmark problems, NPB2.3, was released in 1997.

Since then, high-performance computer systems have grown significantly in processor count and capabilities. Cache and memory capacities have increased, clock rates have gone up, compiler technology has improved, network bandwidths have increased, and total numbers of processors have gone up. Consequently, more challenging benchmark sizes have become necessary to rate the performance of computer systems. At the same time, other important parallelization methods have been standardized, warranting alternate source code implementations.

As parallel computing matured, practitioners also started to pay attention to more practical issues, most notably how to store and process the results of their computations; writing and accessing computed solutions on disk proved a major bottleneck, and ways to assess the efficiency of proposed solutions to this I/O problem became necessary.

Recently, the new field of grid computing has gained popularity. As pioneering efforts are underway to make these grids usable and efficient, NAS researchers have already recognized the need for yardsticks to rate their performance.

To meet these needs, NAS researchers have developed new versions of and extensions to NPB. These have been released in three packages, NPB2.4, NPB3.0, GridNPB3.0, which are available for worldwide distribution and can be downloaded from the NAS Software Repository: www.nas.nasa.gov/Software. More details can be found in NAS technical reports, available on the web at www.nas.nasa.gov/Research/Reports/techreports.html, and efforts are underway to streamline the process for submitting and viewing new results online. For information on NPB, visit www.nas.nasa.gov/Software/NPB.

Editor's Note: *Michael Frumkin, Henry Jin, and Parkson Wong also contributed to this article.*

NAS Grid Benchmarking Software Now Available On Web

NPB2.4: Larger problems, I/O benchmark

A relatively straightforward extension of the NPB MPI implementation, NPB2.4-MPI provides class D, a new set of problem sizes that involves approximately 20 times as much computational work as NPB2.3's class C. Dataset sizes have simultaneously been increased by a factor of 16 to keep up with the larger cache memories. In addition, an I/O benchmark program, previously called BTIO, is now being released under the name NPBIO2.4-MPI.

This version is virtually the same as the BT benchmark program in NPB2.4-MPI, but also writes the computed solution to disk at set intervals, using any of four methods to expedite the process. Users can select the method whose performance they want to rate, one of which is based on high-level, efficient I/O operations defined in the most recent version of the MPI standard. NPB2.4-MPI and NPBIO2.4-MPI are released together in a package called NPB2.4.

NPB3.0: New parallelizations, new language

While MPI was the first widely accepted standard for parallelizing scientific programs, computational scientists have since proposed several other paradigms, which has led to some new NPB implementations. Prototypes of these codes were initially included in a limited distribution, then called Programming Baseline for the NAS Parallel Benchmarks (PBN). For more information on this, see *NAS News*, Summer 1999, on the web at:
www.nas.nasa.gov/About/Gridpoints/PDF/nasnews_V04_N03_1999.pdf.

Now that these new parallelization standards have reached a significant level of acceptance and maturity, NAS researchers have decided to incorporate them into a full, worldwide NPB release. The researchers based their efforts on an improved NPB version without message passing constructs, designated NPB3.0-SER (SER stands for "serial," that is, written in conventional, non-parallel Fortran), and which they are now releasing as a service to researchers who want to test new parallelization tools.

The new parallel versions, called NPB3.0-OMP, NPB3.0-HPF, and NPB3.0-JAV, are based on OpenMP, High Performance Fortran, and Java threads, respectively. Together with NPB3.0-SER, they are being released under the name NPB3.0. OpenMP has become particularly popular on shared-memory systems because of its ease of program development. High Performance Fortran, generally considered a niche programing language in high-performance computing, was recently cited in a significant fluid flow computation on the Japanese Earth Simulator, logging an impressive 12.5 Teraflops (floating point operations per second). Java, though not usually associated with high-performance computing, was included because of its excellent portability properties.

GridNPB3.0: Benchmark for grid computing

Computational grids, encompassing many high-performance systems connected in ultra-wide-area networks, have recently been receiving a lot of attention among research and industry groups. Several solutions have been proposed to solve the problem of the software interface to services on such sprawling systems, with NASA's own effort, the Information Power Grid, an important player. Technical successes have been reported at a number of scientific meetings, but little quantitative data is readily available for use by application engineers interested in the efficiency and ease of use of grid computing. The NAS Grid Benchmarks (NGB), (see *Gridpoints*, Summer 2002, page 5) aim to fill that gap. They are a paper-and-pencil specification of four families of problems that represent important applications on computational grids, such as parameter studies involving many independent runs of the same program, each with a different value of some problem parameter(s).

Another important, slightly more complex problem is that of the so-called visualization pipe, in which a sequence of computational, postprocessing, and visualization tasks is executed cyclically within a time-dependent calculation. Such a process, shown schematically in the figure on page 18, is a typical example of what users may want to run on a grid, and is captured in one of the four benchmark problems.

While it is too early to declare a winner in the race to usable grid software, NAS researchers decided to provide some correct – though not necessarily efficient – reference implementations of NGB to help other groups benchmark their systems. Called GridNPB3.0 because they use NPB3.0 components as building blocks, these implementations consist of Fortran programs executing serially and communicating through a shared file system (GridNPB3.0-SHF-SER), or of Fortran or Java programs executing fully concurrently (both inter- and intra-task parallelism), communicating through the Java Remote Method Invocation (Grid NPB3.0-JAV-OMP/JAV).

NAS Technical Training Seminars

Many of the seminars presented at the NASA Advanced Supercomputing (NAS) Facility are videotaped, and can be borrowed by sending e-mail to the NAS Documentation Center (doc-center@nas.nasa.gov). The seminars highlighted below are all available on videotape. Please reference the date code when ordering.

Microphysics, Radiation, and Surface Processes

Wei-Kuo Tao from NASA Goddard Space Flight Center's Laboratory for Atmospheres spoke about "Microphysics, Radiation and Surface Processes in the Goddard Cumulus Ensemble (GCE) Model" at a New Technology Seminar. The GCE Model has been developed and improved at Goddard over the past two decades, and has been used to understand a variety of climate phenomena. These include: water and energy cycles and their roles in the tropical climate system; the vertical redistribution of ozone and trace constituents over various spatial scales by individual clouds and well-organized convective systems; the relationship between the vertical distribution of latent heating (phase change of water) and the large-scale (pre-storm) environment; and the validity of assumptions used in the representation of cloud microphysical processes and their interaction with radiative forcing over tropical and mid-latitude regions. Tao discussed five major GCE improvements, as well as their performance. (September 17)

Cluster Approaches to Solvation and Surface Chemistry

Mark Gordon from the Department of Chemistry and the Ames Laboratory for Scalable Computing at Iowa State University discussed "Cluster Approaches to Solvation and Surface Chemistry." For the past several years, researchers at Iowa State University have been developing methods to approach condensed phase behavior, starting from a molecular, as opposed to a continuum basis. This has resulted in the effective fragment potential (EFP) model, which has been applied most extensively to problems in which water is the solvent. Gordon discussed an embedded cluster method for treating surface chemistry known as the surface integrated molecular orbital-molecular mechanics (SIMOMM) method, which integrated quantum mechanics with molecular mechanics. (September 12)

Higher Order Godunov Finite-Volume Approximation of PDEs

Tim Barth, from the NASA Advanced Supercomputing Division at Ames Research Center, described his research "Genuine Error Estimates for the Higher Order Godunov Finite-Volume Approximation of Partial Differential Equations (PDE)" at a new technology seminar. Barth presented a-posteriori error estimates for high-order Godunov finite-volume methods which exploit the inherent two-solution representations. He also discussed important issues such as the treatment of nonlinearity and postprocessing of dual data. In addition, Barth presented numerical results for linear and nonlinear conservation laws to verify the analysis. (August 22)

Tool Support for OpenMPI

Barbara Chapman, associate professor of Computer Science at the University of Houston, discussed "Tool Support for the OpenMPI API." Chapman noted that OpenMP is the language of choice for shared memory parallel programming and, having undergone minor revisions within the past two years, is now relatively stable. OpenMP is also strongly supported by the major vendors. Chapman believes that tools are needed to support application developers who use OpenMP, whether they are using it alone or in conjunction with MPI for better portability. At this time, however, relatively few commercial tools are available. Chapman presented a technology to help researchers create and share tools that support the creation and optimization of OpenMPI programs. (July 8)

Computational Nanotechnology of Nanomaterials and Devices

Deepak Srivastava, senior scientist and task lead in the NAS Division's computational nanotechnology group, was the invited speaker at the Information Sciences and Technology Directorate's monthly "IT Forum." Srivastava spoke on "Computational Nanotechnology of Nanomaterials and Devices." The role of computational nanotechnology in advancing the possibilities in nano-materials, electronics, sensors, and machines was described as was the use of carbon nanotubes and fullerenes and how they apply to NASA's missions. Using large-scale classical atomistic and quantum molecular dynamics and transport simulations, Srivastava discussed the nanomechanics of carbon and boron-nitride nanotubes in ultra-light and very strong functional composite materials; branched carbon nanotube structures for biomimetic molecular electronics and sensing applications; endo-fullerenes and doped diamond nanocrystals for solid-state quantum computers; and model operations of molecular gears and motor systems. Large-scale physics and chemistry-based modeling and simulation tools and capabilities have a critical role to play in the development of nanotechnology and new frontiers at the interface of info-bio-nano technologies. Srivastava also discussed near-term requirements, and areas for pushing this growth, within a core information technology-based environment. (September 13)

Calendar of Events

Supercomputing 2002 (SC2002)

Baltimore, Maryland • November 16–22

“From Terabytes to Insights” is the theme for this year’s supercomputing conference. SC2002 will bring together scientists, engineers, designers, and managers from all areas of high-performance networking and computing, showcasing the latest in systems, applications, and services. Check out the conference website at: www.sc2002.org. See the special insert in this issue beginning on page 1A.

3rd International Workshop on Grid Computing

Baltimore, Maryland • November 18

This international meeting is an opportunity for the grid community to discuss current and future work, as well as exchange research ideas in this field. The focus of this meeting will be on grid applications. Find out more at: www.gridcomputing.org/grid2002

41st AIAA Aerospace Sciences Meeting

Reno, Nevada • January 6–9, 2003

The Aerospace Sciences Meeting is the largest of the AIAA technical conferences, and is the preeminent technical gathering for computer and software technologies applicable to the aerospace industry. Technical papers are presented by authors chosen via a competitive selection process based on peer review. Further details are available at: www.aiaa.org

17th Annual HPCC Conference

Newport, Rhode Island • March 25–27, 2003

The National High Performance Computing and Communications Council holds an annual High Performance Computing and Communications Conference each spring.

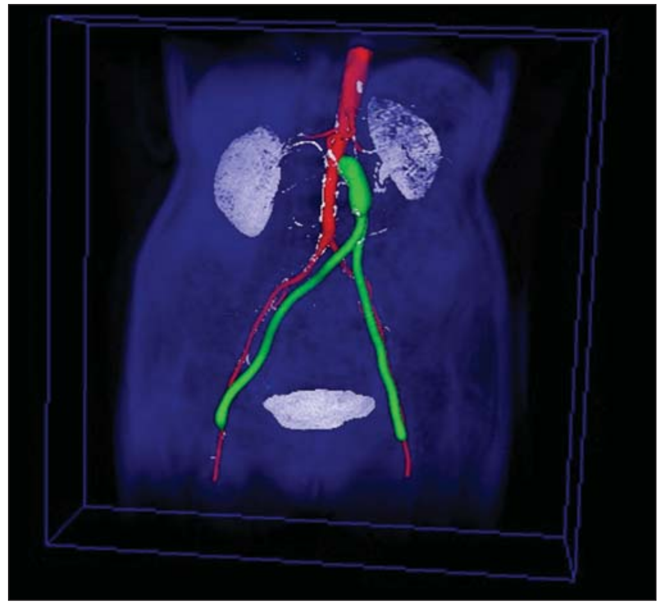

Researchers at Stanford University and NASA are working to create a simulation-based medical planning system for cardiovascular patients that uses computational methods to evaluate alternative surgical options prior to treatment, using patient-specific models of the vascular system. Here a patient-specific preoperative geometric model (shown in red) is created from 3-D preoperative medical imaging data. An operative plan (shown in green) is then created by a surgeon. This technology will be demonstrated at the SC2002 conference in Baltimore, November 16–22. (Stanford University)

This is one of the few conferences that emphasize communication between manufacturers and users, as well as academics and the government agencies which establish policy and regulate the use of advanced technologies. Topics covered include: wireless computing, e-government, Internet security, grid computing, mass storage, homeland cyber security, and bioterrorism. For more information, visit: www.hpcc-usa. org/genconf.ht

2002 NASA Information Power Grid (IPG) Tutorial and Workshop

Palo Alto, California • December 10-12

The IPG Tutorial will be held December 10, providing an introduction to the NASA grid and its various tools. This is an excellent opportunity for people to become familiar with the general capabilities of the IPG, the basic functions available to its users, and those developing IPG applications. The tutorial requires no previous IPG experience.

In contrast to other grid workshops and events, the IPG workshop is focused specifically on NASA's Information Power Grid. The objectives of the IPG workshop are to provide a forum in which: IPG system developers and implementors can share their experiences; IPG application developers and users can share their experiences; potential IPG users can gain a broad understanding of its capabilities; and IPG research and development personnel and the NSF PACI partners can report on the progress of their grid research and development.

Presentations will include talks from research and development teams from each of the major IPG-related organizations including: NASA centers; Argonne National Laboratory; National Computational Science Alliance; San Diego Supercomputer Center; University of Southern California’s Information Sciences Institute, and other participating universities.

For additional information on this event, point your browser to: www.nas.nasa.gov/User/Training/training.html, and click on “Workshops/Conferences,” or contact Marcia Redmond at: mredmond@mail.arc.nasa.gov.

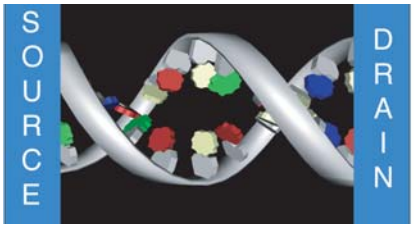

Proposed models of nanoscale transistors

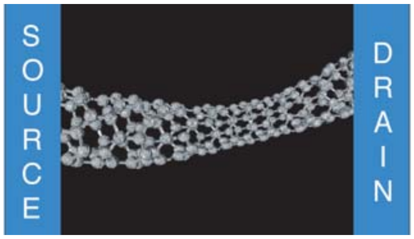

Nanotechnology: Simulating the Future of Computing

NAS scientists demonstrate how the laws of quantum mechanics affect transistor and device miniaturization and construction. See page 8.

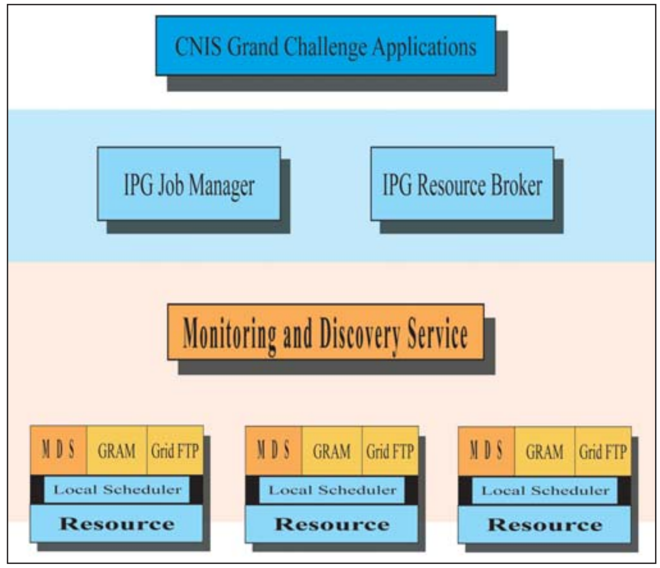

NASA's IPG Team Develops New Tools

NASA's IPG team meets another of its milestones, creating tools to simplify job submission to distributed grid resources. *See page 4.*

www.nas.nasa.gov/gridpoints

National Aeronautics and Space Administration
NASA Advanced Supercomputing Division
Ames Research Center
MS IN258-6
Moffett Field, CA 94035-1000

Official Business
Penalty for Private Use $300

FIRST CLASS MAIL
POSTAGE & FEES PAID
NASA
Permit No. G-27

Address Correction
If the above address or mail stop is incorrect, please e-mail a current address to gridpoints@nas.nasa.gov, or write to:
Gridpoints Editor, MS 258-6, Moffett Field, CA 94035

PLEASE RECYCLE
Printed on recycled and recyclable paper with vegetable-based inks